Granular Instrumental Variables*

Xavier Gabaix and Ralph S.J. Koijen

February 8, 2019. Incomplete. Please do not circulate.

Abstract

In many settings, there is a dearth of instruments, which hampers economists' ability to investigate causal relations. In this paper we propose a quite general way to construct instruments: "granular instrumental variables" (GIVs). In the economies we study, a few large firms or countries account for a large share of economic activity. As they are large, their idiosyncratic shocks affect aggregate outcomes. This makes those idiosyncratic shocks valid instruments for the aggregate shocks. We provide a methodology to extract idiosyncratic shocks from the data, this way creating GIVs. Those GIVs allow us to then estimate parameters of interest, including causal elasticities.

We first illustrate the idea in the basic supply and demand framework: we achieve a novel identification of supply and demand elasticities, based on idiosyncratic shocks to supply or demand. We then show how the procedure can be adapted to handle many enrichments.

We illustrate the procedure in detail with a few applications. First, we measure how shocks to domestic banks causally affect sovereign yields. We document how negative shocks to Italian banks adversely affect Italian government bond yields, and vice-versa. This gives the first causal measure of the "doom loop" between banks and sovereign yields. Second, we study the impact of mega-firms on aggregates, such as aggregate competition (as idiosyncratic shocks to large firms change concentration ratios). This allows us to find a causal relationship between the rise of mega firms and lower investment. Third, we study supply and demand elasticities for oil. This allows us to compare the GIV estimate to previous estimates, this way confirming the validity of our approach.

GIVs allow for a resolution of the still-open question of how to go from "micro-elasticities" to "macro-elasticities". We study a simple model multi-region model with consumption linkages and pricing frictions. We find that, under some plausible conditions, macro-elasticities are the micro elasticities times the "GDP spillover" M (which indicates how much growth in one region affects growth in other regions), which our GIV method estimates to be about 1.5. Hence, the typical micro-elasticities need only to be multiplied by this M to become macro-elasticities.

We conclude that the GIVs are a good candidate to supplement the empirical economists' toolbox.

*xgabaix@fas.harvard.edu, ralph.koijen@chicagobooth.edu. We thank Antonio Coppola, Rishab Guha, Dong Ryeol Lee for very good research assistance. We thank Gary Chamberlain, Mark Gertler, Oleg Itskhoki, Jim Stock and seminar participants at Harvard, NY Fed, the University of Virginia and the NBER (conference on "the rise of mega-firms") for comments. We are grateful to German Gutiérrez and Thomas Philippon for sharing their data with us.

1 Introduction

In many settings, there is a dearth of instruments, which hampers economists' understanding of causal relations (Ramey (2016),Nakamura and Steinsson (2018),Stock and Watson (2016, 2018); Chodorow-Reich (2017)). In this paper we propose a general way to construct instruments: "granular instrumental variables" (GIVs). Those instruments in turn allow to tease apart causal relations in a wide variety of economic contexts.

In the economies we study, many decisions are taken by a few large firms (we use "firms," but soon we will generalize to large entities such as industries or, in an international context, countries). As has been observed before, those firms are often large, and their idiosyncratic shocks (e.g., productivity shocks) affect the aggregate.[1]

Those idiosyncratic firm or country level shocks are valid instruments for the aggregate shocks. Now, we need a methodology to extract those idiosyncratic shock, and the paper presents such a methodology. This creates GIVs. Those GIVs then allow us to estimate parameters of interest.

We first illustrate the idea in a basic static setup with supply and demand (section 2). It is a classic case, and we show how GIVs allow for a novel estimation procedure: they yield an instrument that allows us to estimate the elasticities of supply and demand.[2] Indeed, idiosyncratic demand shock to large firms or countries give a valid instrument for demand change – and thus allow one to estimate the elasticity of supply. They also allow us to estimate the elasticity of demand: the idiosyncratic demand shock of a large firm impacts the price, which changes the demand of other firms — thus allowing to estimate the demand elasticity. We formalize those ideas, and give a way to "optimally extract" idiosyncratic shocks, this way constructing optimal GIVs.[3]

Once the ideas are in place, we show in Section 3 how the procedure can be extended to handle many enrichments, such as feedback loops, heterogeneity, and multidimensionality.[4] We specify the procedure for this general framework.

Empirical illustrations We provide preliminary empirical results for three applications (Section 4). First, we study the link between fragility in the banking sector and sovereign bond markets, so-called "financial doom loops." The European sovereign debt crisis highlighted the possibility

[1] Hence, economies are "granular," that is, and their shocks made of incompressible "grains" of economic shocks, those are firm or industry-level shocks. This theme is laid out in Gabaix (2011), and developed in e.g. Acemoglu et al. (2012); Di Giovanni and Levchenko (2012); Di Giovanni et al. (2014).

[2] The idea that we propose is, in retrospect, so natural that we suspected that it may have been already proposed in the literature. However, after searching the literature and consulting with professional econometricians, we could not find it. Still, we cannot exclude the possibility that it has been proposed somewhere – and we welcome more pointers to the literature. However, we are quite sure that this idea has not been implemented in mainstream economic applications.

[3] We pursue other examples, in particular the "reflection problem" (Manski (1993))— how to measure economic or social influence. This can be handled with GIVs: the idiosyncratic shocks to "big influencers" (e.g. large firms, or perhaps famous people in the networks) affect the aggregate, hence they allow to estimate the social or economic multiplier.

[4] We will add a vector autoregression in a future draft.

that weakness in the banking sector may spill over to the sovereign, and vice versa. However, identifying contagion in a causal way has been challenging thus far. We study Italy during the European sovereign debt crisis and its aftermath. We show that the Italian banking sector, like many financial sectors in other countries, is highly concentrated. Moreover, as is well known, the correlation between banking returns and sovereign yield changes turned negative during the crisis.

We show how to construct GIVs in this setting using bank-level stock returns. Our results identify a causal link between the sovereign and the banking sectors during the sovereign debt crisis, but this effect turns insignificant if we re-estimate the model since December 2014 when the large-scale asset purchase program was announced by the European Central Bank. During this second period, we no longer find that banks causally affect the sovereign, while the OLS estimate (which merely measures co-movement) remains significant.

In terms of magnitudes, we find that a 50% negative return on banking stocks leads to a 75bp increase in the sovereign yield when the doom loop is operative. During the full sample, we estimate a significant feedback of primitive shocks that are amplified with a multiplier 1.45, which remains significant but smaller at 1.15 during the second sample. To the best of our knowledge, this is the first causally estimated measure of doom loops. Those are local effects, that could become larger closer to bankruptcy.

Second, we study the impact of industry concentration on the investment rate. Both in the aggregate and across industries, increased industry concentration is negatively correlated with changes in investment. Using data compiled by Gutiérrez and Philippon (2017a) from 1995 to 2015 covering over 40 industries, we show how to construct GIVs for 5-year changes in industry concentration using different assumptions about the nature of the shocks that affect industry concentration, which may help to pin down the mechanism by which concentration impacts investment.

A first insight is that the resulting instruments are not weak, and can plausibly be strengthened further with more detailed data. Our instruments may therefore be useful in other settings to understand the impact of changes in industry concentration on economic outcomes. Second, the fact that increased concentration negatively impacts investment is robust to common shocks to all industries that may differentially impact small and large firms, or common shocks to manufacturing and non-manufacturing industries (but potentially with a differential impact on small and large firms). The link weakens when we allow for industry-specific shocks that may have a heterogeneous impact on firms across the size distribution.

Going from micro to macro elasticities Third, we study GDP spillovers. If Germany enters in a recession, how does that impact the rest of the Eurozone, or the world? To start attacking this questions, we consider US states. If California has a high growth rate, what's the spillover on the other US states? We find a multiplier around 1.5 – albeit with fairly large standard errors (around 0.5). This means that the "direct" impact of California on the US growth rate is multiplied (once GE effects are taken into account) by a factor of $M \simeq 1.5$.

GIVs allows for a resolution of the still-open question of how to go from "micro-elasticities" to "macro-elasticities" (see Section 5). We find theoretically that macro-elasticities are the micro elasticities times the "GDP spillover" M (how much growth in one region affects growth in other regions), which our GIV method estimates to be about 1.5. Hence, the typical "micro-elasticities" need only to be multiplied by this M to become macro-elasticities. For instance, call ψ^ℓ the elasticity of local employment to local housing prices, as in Mian and Sufi (2014). Then, the aggregate impact (i.e. the elasticity of total employment to an aggregate change in housing prices) is: $\psi = M\psi^\ell$.

We think that this methodology will be useful in other contexts, for instance, to study the impact of bank capitalization on lending, the impact on uncertainty on aggregate outcomes (Bloom (2009)), the reflection problem in social science (Manski (1993)), the impact of investors' demand on asset prices (using the methodology in Koijen and Yogo (forthcoming)), or international trade (Di Giovanni et al. (2014); Gaubert and Itskhoki (2018)). We are working on some of those specific examples, and they will be included in this paper as it progresses.

Related literature We relate to a number of literatures.

Instruments for macro. It's active literature (Ramey (2016), Nakamura and Steinsson (2018), Stock and Watson (2018), Chodorow-Reich (2017)). We add to it, by proposing to use GIVs, which are quite ubiquitous. A productive literature studies the impact of the entry of China in the WTO to study various outcomes, the "China shock" (Autor et al. (2013); Acemoglu et al. (2016)). In some sense, GIVs are a systematic generalizations the China shock. There are lots idiosyncratic shocks, and GIVs allows to construct them systematically.

Origins of aggregate fluctuations. We provide tools that can tease that apart when you have loops. This literature is finding that a sizable amount of volatility is "granular" in nature – coming from idiosyncratic shocks to firms or industries (see Long and Plosser (1983); Gabaix (2011), Acemoglu et al. (2012); Di Giovanni and Levchenko (2012); Di Giovanni et al. (2014); Baqaee and Farhi (2018); di Giovanni et al. (2018)). Those datasets can be revisited forming GIVs which allow to investigate causal relations. Gabaix (2011) introduces the notion of "granular residual" – a weighted sum of proxies for idiosyncratic shocks, and shows how idiosyncratic shocks to firms appear to explain about 1/3 of GDP fluctuations. But this paper does not take the crucial step to use this kind of concept as an instrument to measure causal relations, e.g. in a demand and supply setting.

Leary and Roberts (2014) use the actions of competitors as an instrument, but they use an equal-weighted average, so that they do not deploy the power of using the value-weighted minus equal-weighted actions. Amiti et al. (2016) study strategic complementarities in price-setting across Belgian manufacturing firms; as in some of our examples, they seek to identify the ways in which a shock to one firm propagates through the competitive network. They achieve identification through a micro-level analysis of firm exposure to international shocks. Our approach is much more general, and does not require the same level of detailed micro data. The use of model-based idiosyncratic

shocks has some similarities with Amiti and Weinstein (forthcoming), who extract bank supply shocks from Japanese data using a panel of fixed effects, and then estimate the sensitivity of aggregate investment to these shocks. However, unlike our model, Amiti and Weinstein (forthcoming) do not allow for "reflection": shocks to bank affect aggregate investment, but aggregate investment does not circle back around to affect individual bank behavior. This is the key source of endogeneity in many of the models we consider, and by tackling it we are able to estimate a richer set of parameters.

Concentration and competition. A number of authors have established correlations between changes in concentration and various economic outcomes like the labor share (Autor et al. (2017)). Here the GIV is made of firm-specific growth shocks, that shock the concentration (e.g., Herfindahl or other concentration ratio) of the economy. GIVs are available for all periods, while other instruments (e.g. China shock) are only available for special time periods and industries. Hence, GIVs complement the studies of investment (Gutiérrez and Philippon (2017a); Crouzet and Eberly (2018)) and wages (Azar et al. (2017)).

We also relate to more methodological literatures.

Reflection problem. There is a rich literature studying social interactions (Manski (1993)). Interestingly they eliminate any use of the noise. We identify things where it seems previously hopeless (Sections 2.4 and 10.7) – this literature hasn't explored systematically using idiosyncratic shocks.

Spatial econometrics. In part of our work, growth in a region affect the other regions. So, in that part there is a similarity between our setup and that of spatial econometrics (e.g. Kelejian and Prucha (1999); Blasques et al. (2016); Shi and Lee (2017); Kuersteiner and Prucha (2018)). However, the estimators are quite different. The reason is that spatial econometrics studies the "local" influence (e.g. of neighboring cities on a city), while GIVs study the global influence. Hence, the sources of variations, identifiability conditions and methods are quite different. Certainly, the spatial literature has not identified, as we do, the GIVs as a simple way to estimate elasticities in contexts such as supply and demand problems, and models with feedback loops e.g. from banks to sovereign yields.

The reader is encouraged to read Section 2 and 4 at first, and then go to Section 3, which contains the general procedure, in a second reading.

2 Granular identification: The basics

2.1 Notations

Notations For a vector $X \coloneqq (X_i)_{i=1...N}$ and a series of time-invariant relative weights S_i with $\sum_i S_i = 1$, we call

$$X_E \coloneqq \frac{1}{N}\sum_{i=1}^{N} X_i, \tag{1}$$

$$X_S \coloneqq \sum_{i=1}^{N} S_i X_i, \tag{2}$$

$$X_\Gamma \coloneqq X_S - X_E, \tag{3}$$

so that X_E is the equal weighted average, X_S is the size weighted average.

In the paper, we shall also have shocks u_i, uncorrelated with variance $\sigma^2_{u_i}$. Then, we'll define the "pseudo-equal weight" or "inverse variance weights"

$$\tilde{E}_i \coloneqq \frac{1/\sigma^2_{u_i}}{\sum_j 1/\sigma^2_{u_j}}, \tag{4}$$

which satisfies $\sum_i \tilde{E}_i = 1$, and is equal to $\tilde{E}_i = \frac{1}{N}$ in when all the $\sigma^2_{u_i}$ are equal. We'll also define

$$\tilde{\Gamma}_i \coloneqq S_i - \tilde{E}_i. \tag{5}$$

Then, $X_{\tilde{\Gamma}}$ will be the "granular residual" in a number of settings. It is the size-weighted average of X minus the "inverse-variance" weighted average of X. It will be an optimal proxy for idiosyncratic shocks.

We use notation β^e for the estimator of a parameter β; $\mathbb{E}_T$ for the sample temporal mean, $\mathbb{E}_T[Y_t] \coloneqq \frac{1}{T}\sum_{t=1}^{T} Y_t$; C_t for a vector of controls; ι for a vector of 1's, I for the identity matrix, of the appropriate dimension given the context; $V^Y = \mathbb{E}[Y_t Y_t']$ for a variance-covariance matrix.

Herfindahl We define the Herfindahl (strictly speaking, this is an "excess Herfindahl") as:

$$h \coloneqq \sqrt{-\frac{1}{N} + \sum_{i=1}^{N} S_i^2}. \tag{6}$$

A higher $h \in \left[0, \sqrt{1-\frac{1}{N}}\right]$ means that the industry is more concentrated: an industry where all the firms have the same size features $h = 0$.

Variance facts We will repeatedly use the fact that if $(u_i)_{i=1...N}$ is a series of uncorrelated random variables with mean 0 and common variance σ_u^2, then

$$\mathbb{E}\left[u_\Gamma u_E\right] = 0, \tag{7}$$

and

$$\mathbb{E}\left[u_\Gamma^2\right] = \mathbb{E}\left[u_S u_\Gamma\right] = h^2\sigma_u^2. \tag{8}$$

Hence, the standard deviation of the granular residual $u_{\Gamma t}$ is proportional to the Herfindahl. In the general heteroskedastic case, we have also (the proof is in Section 19.9)[5]

$$\mathbb{E}\left[u_{\tilde{\Gamma}} u_{\tilde{E}}\right] = 0. \tag{9}$$

2.2 A very simple example with no feedback loop

We consider a very simple example with supply and demand.

2.2.1 Basic model

For clarity, we lay out a concrete economic models. Demand by country i at date t is $D_{it} = \bar{Q}S_i\left(1 + y_{it}\right)$, where $\bar{Q}$ is the average total world production, and S_i is country $i's$ share of demand, normalized to follow:

$$\sum_{i=1}^{N} S_i = 1. \tag{10}$$

The demand disturbance is assumed to be:

$$y_{it} = \eta_t + u_{it}. \tag{11}$$

It is the sum of a common shock η_t, and an idiosyncratic shock u_{it}. All shocks are IID across dates. Then, total world demand is $D_t = \sum_i D_{it} = \bar{Q}\left(1 + y_{St}\right)$, where $y_{St} \coloneqq \sum_i S_i y_{it}$ is the size-weighted average demand disturbance. We suppose that supply is $Q_t = \bar{Q}\left(1 + \frac{p_t - \varepsilon_t}{\alpha}\right)$, where $p_t = \frac{P_t - \bar{P}}{\bar{P}}$ is the proportional deviation from $\bar{P}$, which is thus the average price of oil. Then, to equilibrate supply and demand ($D_t = Q_t\left(p_t\right)$), we must satisfy: $\bar{Q}\left(1 + y_{St}\right) = \bar{Q}\left(1 + \frac{p_t - \varepsilon_t}{\alpha}\right)$, i.e. the deviation of the price from the average satisfies:

$$p_t = \alpha y_{St} + \varepsilon_t. \tag{12}$$

[5]Here is the proof. We have $E = k\left(V^u\right)^{-1}\iota$ for $k = \frac{1}{\iota\left(V^u\right)^{-1}\iota}$. So

$$\mathbb{E}\left[u_{\tilde{\Gamma}} u_{\tilde{E}}\right] = \mathbb{E}\left[\left(E'u\right)\left(u'\Gamma\right)\right] = E'\mathbb{E}\left[uu'\right]\Gamma = E'V^u\Gamma = k\iota'\left(V^u\right)^{-1}V^u\Gamma = k\iota'\Gamma = 0$$

as $\iota'\Gamma = 0$.

It depends on the size-weighted average demand shock. This is, the outcome p_t depends on the size-weighted average demand disturbances, $y_{St} = \sum_i S_i y_{it}$. The aggregate demand and supply shocks η_t, ε_t can be correlated.

The classic problem is that we cannot estimate α by OLS. Indeed, a regression of p_t on y_{St} ($p_t = \alpha y_{St} + \varepsilon_t$) would be biased, as ε_t and y_{St} are correlated.

However, suppose that we form the GIV:

$$z_t \coloneqq y_{\Gamma t} = y_{St} - y_{Et} = \sum_{i=1}^{N} S_i y_{it} - \sum_{i=1}^{N} \frac{1}{N} y_{it}. \tag{13}$$

Then, we have (using $u_{St} \coloneqq \sum_{i=1}^{N} S_i u_{it}$, $u_{Et} \coloneqq \sum_{i=1}^{N} \frac{1}{N} u_{it}$)

$$y_{St} = \sum_{i=1}^{N} S_i y_{it} = \eta_t + u_{it}, \qquad y_{Et} = \sum_{i=1}^{N} \frac{1}{N} y_{it} = \eta_t + u_{Et},$$

so

$$z_t = (\eta_t + u_{S_t}) - (\eta_t + u_{Et}) = u_{St} - u_{Et} =: u_{\Gamma t}.$$

Note that $z_t = y_{St} - y_{Et}$ is just formed by observables. It is the difference between the size-weighted demand and the equal weighted demand. Intuitively, it captures the "idiosyncratic demand" by large units, as shown by $z_t = u_{\Gamma t}$.

We assume that the shocks u_{it} are idiosyncratic, in the sense that:

$$\mathbb{E}\left[u_{it}\varepsilon_t\right] = 0 \text{ for all } i, t, \tag{14}$$

but the u_{it}, u_{jt} could be correlated. But, in this simple example, all shocks are IID over time.

Then, we have

$$\mathbb{E}\left[z_t \varepsilon_t\right] = 0: \text{ Exogeneity}, \tag{15}$$

and

$$\mathbb{E}\left[z_t y_{St}\right] \neq 0: \text{ Relevance}. \tag{16}$$

Hence, $z_t = u_{\Gamma t}$ is a valid instrument. We call it a "granular instrumental variable" (GIVs).

Given that $p_t - \alpha y_{St} = \varepsilon_t$, we have

$$\mathbb{E}\left[\left(p_t - \alpha y_{St}\right) z_t\right] = 0, \tag{17}$$

i.e. $\mathbb{E}\left[p_t z_t\right] - \alpha \mathbb{E}\left[y_{St} z_t\right] = 0$, which gives the supply elasticity α, by

$$\alpha = \frac{\mathbb{E}\left[p_t z_t\right]}{\mathbb{E}\left[y_{St} z_t\right]}. \tag{18}$$

Indeed, in practice, we might estimate α by using the temporal means:

$$\hat{\alpha}_T \coloneqq \frac{\frac{1}{T}\sum_t p_t z_t}{\frac{1}{T}\sum_t y_{St} z_t} = \frac{\mathbb{E}_T\left[p_t z_t\right]}{\mathbb{E}_T\left[y_{St} z_t\right]}. \tag{19}$$

We state a formal proposition.

Proposition 1 (Consistency of the GIV estimator in this example). *Suppose that only that* $\mathbb{E}\left[u_{it}\varepsilon_t\right] = 0$ *(but the* u_{it} *can have an arbitrary distribution, with mean 0), and a succession of i.i.d. dates* t*. Form the GIV estimator* $z_t \coloneqq y_{\Gamma t}$*. Then,* z_t *identifies the price elasticity, by* $\alpha = \frac{\mathbb{E}[p_t z_t]}{\mathbb{E}[y_{St} z_t]}$*. In other terms, the GIV estimators for the price elasticity* α*,* $\hat{\alpha}_T \coloneqq \frac{\frac{1}{T}\sum_t p_t z_t}{\frac{1}{T}\sum_t y_{St} z_t} = \frac{\mathbb{E}_T[p_t z_t]}{\mathbb{E}_T[y_{St} z_t]}$*, is a consistent estimator.*

OLS estimator version One can also use a plain OLS estimator with our instrument $u_{\Gamma t}$.[6] Assume for simplicity that the u_{it}'s are uncorrelated and have the same variance (Section 19.9 shows how to relax that). Given that

$$p_t = \alpha u_{St} + \alpha\eta_t + \varepsilon_t,$$

we can simply estimate α by OLS:

$$p_t = \alpha z_t + \varepsilon_t^0, \tag{20}$$

with $\varepsilon_t^0 \coloneqq \alpha\eta_t + \varepsilon_t + \alpha u_{Et}$ (indeed, we have $\mathbb{E}\left[u_{\Gamma t} u_{Et}\right] = 0$). Call $\alpha^{G,OLS0}$ that estimator.

We can improve the efficiency of the OLS estimator. We assume $\mathbb{E}\left[u_{it}\eta_t\right] = 0$ for all i and t. We use

$$\eta_t^e = y_{Et}$$

as an estimator of η_t. We have $\eta_t^e = \eta_t + u_{Et}$, and $u_{Et} = O\left(\frac{1}{\sqrt{N}}\right)$ can be expected to be small. Hence, regress:

$$p_t = \hat{\alpha} u_{\Gamma t} + \hat{\beta}\eta_t^e + \varepsilon_t^p, \tag{21}$$

which will yield in the limit of large T a consistent estimate of α. [7] In addition, it turns out that it is a more precise estimator (when N is large enough), as we control for η_t.

Precision of GIV estimators The next proposition states when we have a precise estimator (its proof is in Section 12).

[6] In estimator 19, the denominator could be of either sign, and close to 0 in finite sample, leading to some instability in the estimate.

[7] However, we will discard $\hat{\beta}$ as it does not estimate β : it is polluted by the correlation with measurement error. Indeed, $\beta = \alpha + b + O\left(\frac{1}{N}\right)$ with a bias $b = \frac{\mathbb{E}[\varepsilon_t \eta_t]}{\mathbb{E}[\eta_t^2]}$

Proposition 2 (Precision of the GIV estimators in this example). *We assume that the u_{it} are uncorrelated but possibly heteroskedastic. The above estimators based on the granular instrument variable (GIV) $y_{\Gamma t}$ all achieve identification of the reaction parameter α. The GIV and two G-OLS estimators have standard errors*

$$\sqrt{T}\left(\hat{\alpha}_T^q - \alpha\right) \sim \mathcal{N}\left(0, \sigma_{\alpha^q}^2\right),$$

where the standard error of estimator q (equal to GIV, G-OLS0 or G-OLS1) has the form

$$\sigma_{\alpha^q} = \frac{\sigma_{\varepsilon^{(q)}}}{\sigma_{u_\Gamma}}, \tag{22}$$

with[8]

$$\sigma^2_{\varepsilon^{(G-OLS0)}} = var\left(\varepsilon_t + \alpha\eta_t + \alpha u_{Et}\right), \qquad \sigma^2_{\varepsilon^{(IV)}} = var\left(\varepsilon_t\right), \qquad \sigma^2_{\varepsilon^{(G-OLS1)}} \coloneqq var\left(\varepsilon_t^{\perp}\right).$$

We always have $\sigma_{\alpha^{G-OLS}} > \sigma_{\alpha^{IV}}$, and, if $corr\left(\varepsilon_t, \eta_t\right) \neq 0$ and N is large enough, we have $\sigma_{\alpha^{IV}} > \sigma_{\alpha^{G-OLS0}}$.

When the u_{it}'s are IID, we have

$$\sigma_{u_\Gamma} = h\sigma_u. \tag{23}$$

So to have a precise estimate (low σ_α) we need: a few large firms (to have a large excess Herfindahl h), and that demand shocks do affect the price importantly, compared to aggregate shocks (large $\sigma_u/\sigma_\varepsilon$).

This simple example illustrates the basics. In the full version, we can handle a much richer setup, including with feedback loop, richer heterogeneity, autocorrelations (and indeed, vector autocorrelation (VAR)), and multidimensionality. We next move to a simple example with feedback loops.

2.2.2 Model with an enriched factor structure

The model might have a richer factor structure, with r factors, i.e.

$$y_{it} = \sum_{f=1}^{r} \lambda_i^f \eta_t^f + u_{it}, \tag{24}$$

or, in vector form:

$$y_t = \Lambda\eta_t + u_t. \tag{25}$$

where Λ is a $N \times r$ matrix.

[8] Here $\varepsilon_t^{\perp}$ is the residual from projecting $\varepsilon_t - \alpha u_{Et}$ on $\eta_t + u_{Et}$. It is detailed in the proof.

Then, we simply run a factor model (i.e. via Principal Component Analysis (PCA)), and, essentially, we extract the residuals u_{it} to form the GIV. Let us see that more precisely. Suppose that we have estimated the λ vector (e.g. via PCA, as we will detail later). Then, call Q a $N \times N$ matrix projecting vectors on a space orthogonal to Λ, so that $Q\Lambda = 0$ (later Q will be optimally chosen). Then,

$$Qy_t = Qu_t.$$

Then, the GIV is formed as:

$$z_t \coloneqq S'Qy_t, \tag{26}$$

so that

$$z_t = S'Qu_t. \tag{27}$$

Then, z_t is a valid instrument: as $p_t - \alpha y_{St} = \varepsilon_t$ and $\mathbb{E}[u_t \varepsilon_t] = 0$, we have (17), $\mathbb{E}[(p_t - \alpha y_{St}) z_t] = 0$. One can also show that this GIV is the optimal one in some sense.[9] It can also be written:

$$z_t \coloneqq \Gamma' y_t = \Gamma' u_t \tag{28}$$

for

$$\Gamma \coloneqq Q'S. \tag{29}$$

In our basic example (11), $Q = I - \iota E'$, so that the GIV was $z_t \coloneqq \Gamma' y_t$ with

$$\Gamma = Q'S = S - E. \tag{30}$$

A concrete way to see the same thing is that extract the residuals from the factor analysis, which will be $\check{u}_t = Qu_t$, and then we form

$$z_t = S'\check{u}_t = \sum_i S_i \check{u}_{it}. \tag{31}$$

In our basic example, $\check{u}_{it} = u_{it} - u_{Et}$, so $z_t = \sum_i S_i \check{u}_{it} = u_{St} - u_{Et}$.

This shows that the GIV is valid and possible as long as

$$\Gamma \coloneqq Q'S \neq 0.$$

Fortunately, this is generically true. If Γ was close to 0, it would be picked up by very large standard errors.

[9] We want $z_t = a'Qu_t$ for some vector a, as z_t can only use the purified residual Qu_t, and be close to $S'u_t$. So it's intuitively clear that we might use (26). More precisely, we want $\max_{a \in \mathbb{R}^N, Q: Q\lambda = 0} corr(z_t, u_{St})$, and this gives the above estimator, with Q in (94).

2.3 A simple demand and supply example with feedback loops

2.3.1 A simple model

We next enrich the previous example, and consider a simple supply and demand example that features a "loop". Suppose that demand for some commodity (say, oil) is:

$$y_{it} = \lambda p_t + \eta_t + u_{it}, \tag{32}$$

and supply is

$$s_t = \mu p_t + \varepsilon_t, \tag{33}$$

where η_t, ε_t can be correlated. We can expect that the demand and supply elasticities (respectively λ and μ) satisfy $\lambda < 0 < \mu$. Again, to be more formal, y_{it}, s_t and p_t are understood are percent deviation of the average demand of country i, supply and price, respectively.[10]

In equilibrium, supply equals demand,

$$y_{St} = s_t,$$

which gives the price

$$p_t = \frac{u_{St} + \eta_t - \varepsilon_t}{\mu - \lambda}. \tag{34}$$

There is a "loop" because demand shocks η_t and u_{it} feed into the price p_t, which then in turns affects demand. The equilibrium quantity produced is

$$s_t = y_{St} = \frac{\mu u_{St} + \mu \eta_t + \lambda \varepsilon_t}{\mu - \lambda}. \tag{35}$$

The classic problem of estimating supply and demand equilibrium quantity s_t and price p_t is that we cannot regress: $s_t = b^e p_t + \varepsilon_t^e$, as ε_t^e and p_t are correlated.

However, suppose that we form the GIV, as in (13)

$$z_t \coloneqq y_{\Gamma t} = \sum_{i=1}^{N} \left(S_i - \frac{1}{N} \right) y_{it}. \tag{36}$$

Then, we have

$$z_t = u_{\Gamma t}. \tag{37}$$

As in the previous example, we assume that the shocks u_{it} are idiosyncratic:

$$\mathbb{E}\left[u_{it}\eta_t\right] = \mathbb{E}\left[u_{it}\varepsilon_t\right] = 0 \text{ for all } i, t. \tag{38}$$

[10] We take the model of Section 2.2, and simply set $y_{it} = \lambda p_t + \eta_t + u_{it}$, where $p_t = \frac{P_t - P_*}{P_*}$ is the proportional deviation from the average.

Then, we have again

$$\mathbb{E}\left[z_t\varepsilon_t\right] = \mathbb{E}\left[z_t\eta_t\right] = 0: \text{ Exogeneity},$$

$$\mathbb{E}\left[z_t p_t\right] \neq 0: \text{ Relevance}.$$

Estimations of supply and demand elasticities by GIV The supply equation (33) implies

$$\mathbb{E}\left[\left(s_t - \mu p_t\right) z_t\right] = 0, \tag{39}$$

which gives the supply elasticity μ, by

$$\mu = \frac{\mathbb{E}\left[s_t z_t\right]}{\mathbb{E}\left[p_t z_t\right]}. \tag{40}$$

Indeed, in practice, we form the temporal average $\hat{\mu}_T = \frac{\mathbb{E}_T[s_t z_t]}{\mathbb{E}_T[p_t z_t]}$.

Now, we want to estimate demand. For that, we make a stronger assumption: we assume that the shocks u_{it} are uncorrelated (we will relax this later). Then, this implies[11]

$$\mathbb{E}\left[u_{Et}u_{\Gamma t}\right] = 0. \tag{41}$$

So, given this, we have: $y_{Et} - \lambda p_t = \eta_t + u_{Et}$, and $\mathbb{E}\left[\left(y_{Et} - \lambda p_t\right) z_t\right] = 0$. This gives an estimate of the demand elasticity λ,

$$\lambda = \frac{\mathbb{E}\left[y_{Et} z_t\right]}{\mathbb{E}\left[p_t z_t\right]}, \tag{42}$$

and the estimator $\hat{\lambda}_T = \frac{\mathbb{E}_T[y_{Et} z_t]}{\mathbb{E}_T[p_t z_t]}$.

Hence, we have measured the supply and demand elasticities.The next proposition records the precision of the estimators.

Proposition 3 (Precision of the GIV estimate in the supply-and-demand example). *The GIV estimators for the price elasticities of supply ($\mu > 0$) and demand ($\lambda < 0$) have the standard error:*

$$\sqrt{T}\left(\hat{\mu}_T - \mu\right) \sim \mathcal{N}\left(0, \sigma_\mu\right), \qquad \sigma_\mu = \frac{\left(\mu - \lambda\right)\sigma_\varepsilon}{h\sigma_u}, \tag{43}$$

$$\sqrt{T}\left(\hat{\lambda}_T - \lambda\right) \sim \mathcal{N}\left(0, \sigma_\lambda\right), \qquad \sigma_\lambda = \frac{\left(\mu - \lambda\right)\sigma_{\eta + u_E}}{h\sigma_u}. \tag{44}$$

As in Proposition 2, we need a high excess Herfindhal to achieve a high precision.

[11] Indeed, in the homoskedastic case, we have

$$\mathbb{E}\left[u_{Et}u_{\Gamma t}\right] = \mathbb{E}\left[\left(\sum_i \frac{1}{N} u_{it}\right)\left(\sum_i \Gamma_i u_{it}\right)\right] = \frac{1}{N}\sum \Gamma_i var\left(u_{it}\right) = \frac{\sigma_u^2}{N}\sum_i \Gamma_i = 0.$$

as $\sum_i \Gamma_i = 0$.

Relaxing the assumptions We can relax the assumptions to the case where the u_{it} are heteroskedastic with variance $\sigma^2_{u_i}$, but are still uncorrelated (Section 19.9 considers the case where the u_{it}'s are correlated).

Estimation by OLS If we want to estimate things by OLS (as we will often do in the general case), we observe that we can rewrite the solutions as:

$$p_t = \frac{1}{\mu - \lambda} u_{St} + \varepsilon^p_t, \qquad s_t = y_{St} = \frac{\mu}{\mu - \lambda} u_{St} + \varepsilon^s_t,$$

where the $\varepsilon^p_t, \varepsilon^s_t$ are linear combinations of ε_t, η_t. So, if we run the OLS, with $z_t = u_{\Gamma t}$,

$$p_t = b^p z_t + \varepsilon^p_t, \tag{45}$$

we estimate

$$b^p = \frac{1}{\mu - \lambda}, \tag{46}$$

which is the sensitivity of the price to supply or demand shock. If we run the OLS regression:

$$y_{St} = M z_t + \varepsilon^s_t, \tag{47}$$

we estimate

$$M = \frac{\mu}{\mu - \lambda}. \tag{48}$$

From M and b^p, we can recover the elasticity μ and λ. This is mathematically exactly the same estimate as the IV estimator. (40), (42).[12]

Standard errors: when "weak instruments" are or are not a problem When estimating via OLS (e.g. b^p and M), the standard errors are reliably estimated by the usual OLS method, even in small samples. When a ratio is performed (e.g. to estimate λ, μ), the usual delta method, or the GMM based on (39), also give standard errors.[13]

Correspondence with Applied Microeconomics: First and Second Stage in IV Let us recast our GIV in the language of applied microeconomics. One can view the first stage as regression of the price on the GIV:

$$p_t = b^p z_t + \beta^p \eta_t + \varepsilon^p_t, \tag{49}$$

[12] Indeed, the OLS estimators are $\hat{M}_T = \frac{\mathbb{E}_T[y_{St} z_t]}{\mathbb{E}_T[z_t^2]}$ and $\hat{\psi}_T = \frac{\mathbb{E}_T[p_t z_t]}{\mathbb{E}_T[z_t^2]}$. We have $\hat{\mu}_T = \frac{\hat{M}_T}{\hat{\psi}_T} = \frac{\mathbb{E}_T[y_{St} z_t]}{\mathbb{E}_T[p_t z_t]}$, which is the same as (40), as $s_t = y_{St}$ in equilibrium.

[13] Because of a potential division by 0, they can go wrong in small samples – then the literature on weak instruments (Stock and Yogo (2002)) might help.

which gives $b^p = \frac{1}{\mu - \lambda}$. The second stage is running supply on the instrumented change in the price $b^p z_t$:

$$s_t = \mu \left(b^p p_t \right) + \beta^s \eta_t + \varepsilon_t^s, \tag{50}$$

which gives μ. Alternatively, one can run the "reduced form equation"

$$s_t = b^s z_t + \beta^s \eta_t + \varepsilon_t^s, \tag{51}$$

which gives a coefficient $b^s = M = \frac{\mu}{\mu - \lambda}$. The supply elasticity is given by:

$$\mu = \frac{b^s}{b^p}, \tag{52}$$

which has the form of a Wald estimator.

The demand elasticity is similar. We can run equal-weighted demand on the instrumented change in the price, $b^p z_t$

$$y_{Et} = \lambda \left(b^p z_t \right) + \beta^y \eta_t + \varepsilon_t^y, \tag{53}$$

which gives the demand elasticity λ.[14] Alternatively, can we run the reduced form equation $y_{Et} = b^y z_t + \beta^y \eta_t + \varepsilon_t^y$ which gives $b^y = \lambda M$, and the demand elasticity is $\lambda = \frac{b^y}{b^p}$.[15]

2.3.2 A richer model with factor structure

We can also have a richer factor structure, as in Section 2.2.2. The demand could be:

$$y_{it} = \lambda_i^p p_t + \sum_f \lambda_i^p \eta_t^f + u_{it}. \tag{56}$$

First, if we know $\lambda_i^p = \lambda^p$ is a common elasticity, we consider $\check{y}_{it} = y_{it} - y_{Et}$. More generally, if we know that λ^p lives in a given vector space Λ^p (i.e., $\lambda^p = \Lambda^p K$ for some K and precified Λ^p), we

[14] Here we used 41, which makes the OLS valid.

[15] In the more general framework, the notion of what is a first vs. second stage is a bit arbitrary. Let us see how to estimate α^f in that first and second stage language. The *first stage* is the regression:

$$y_{St} = b z_t + \beta_\eta^{y_S} \eta_t^e + \beta_C^{y_S} C_{St} + \varepsilon_t^{y_S}, \tag{54}$$

where we regress on z_t, using our recovered factors η_t as controls. From the model, we know that regression coefficient identifies $b = M$, the multiplier.

The *second stage* is the regression:

$$F_t^f = \alpha^f \left(b^e z_t \right) + \beta^{F^f} \eta_t^e + \varepsilon_t^{F^f}, \tag{55}$$

which gives the estimator for α^f. Alternatively, we can regress F_t^f on z_t (with controls) and the coefficient is $\alpha^f M$. When estimating the influence coefficient, we can also view the second stage as doing

$$y_{Et} = \gamma M^e z_t + \beta_\eta^{y_S} \eta_t^e + \beta_C^{y_S} C_{St} + \varepsilon_t^{y_S},$$

which gives γ.

use $\check{y}_t = Q^{\Lambda^p} y_t$. We run a factor model on the $\check{y}_t$, get the residual $\check{u}_t = Q u_t$, and use that GIV.

2.3.3 Refining the exclusion restriction

Now we can refine the "exclusion restriction". So far we assumed that $\mathbb{E}[u_{it}\varepsilon_t] = 0$, i.e. no covariance between idiosyncratic demand and supply shock. If that's not the case, we can also decompose each supply with a factor model:

$$s_{it} = \mu_i^p p_t + \sum_f \lambda_i^{s,p} \eta_t^{s,f} + u_{it}^s. \tag{57}$$

Then, if the US has a "fracking shock" that affects both supply and demand, it will be captured by both u_{it}^s and u_{it}^d for $i = \text{USA}$, but we still maintain $\mathbb{E}\left[u_{it}^x \eta_t^{x'}\right] = 0$ for x, x' in s, d. But now we keep $\mathbb{E}\left[u_{it}^s u_{it}^d\right]$ potentially nonzero.

Then, we have two instruments: the old one, i.e. the demand-based GIV, $z_t^d = u_{\Gamma t}^d$ (replacing putting a d on the demand side), the new supply-based GIV:

$$z_t^s := \Gamma^{s\prime} y_t^s = u_{\Gamma^s t}^s, \tag{58}$$

where $\Gamma^s = S^s - E^s$ is the relative size on the supply side. We can also form the difference:

$$z_t^{d-s} = z_t^d - z_t^s. \tag{59}$$

Then, the "demand minus supply" GIV is always valid, e.g. it satisfies $\mathbb{E}\left[(y_{Et} - \lambda p_t)\, z_t^{d-s}\right] = 0$. If supply and demand shocks are uncorrelated across countries, then both z_t^d and z_t^s are valid instruments.

2.4 Influence and the "reflection problem"

We finish by another example, known as "contagion" or the "reflection problem" (Manski (1993); Ioannides (2013); Kline and Tamer (Forthcoming)). Suppose that actions follow:

$$y_{it} = \gamma y_{St} + \eta_t + u_{it}, \tag{60}$$

where η_t is uncorrelated with the u_{it}. This equation means that y_{it} is influenced by the aggregate action of other agents (γy_{St}), and in addition to the usual aggregate shocks η_t, and idiosyncratic shocks u_{it} (which we assume to be uncorrelated to η_t). The "influence" or "contagion" parameter γ is of high interest.

The GIV approach works as follows. Taking the size-weighted average of (60), we have $y_{St} = \gamma y_{St} + \eta_t + u_{St}$, so that

$$y_{St} = M(\eta_t + u_{St}), \qquad M = \frac{1}{1-\gamma}. \tag{61}$$

We form $z_t \coloneqq y_{\Gamma t}$,which (by (60)) will give $y_{\Gamma t} = u_{\Gamma t}$. Hence, if we regress by OLS:

$$y_{St} = M^e y_{\Gamma t} + \varepsilon_t^y$$

we will have a consistent estimator M^e of the multiplier M – hence of γ.[16] We have a simple GIV approach to the "reflection problem". To the best of our knowledge, this approach is new. Indeed, it may seem to contradict earlier impossibility results. Section 10.7 solves the apparent contradiction. The short summary is that Manski (1993) and Bramoullé et al. (2009) do not consider anything like a GIV, as they immediately reason on averages based on observables, eschewing any exploration of the noise. In contrast, GIVs are all about exploring some structure in the noise — the idiosyncratic shocks of large entities.

2.5 Optimality of the GIV

Once we see that the GIV instrument $y_{\Gamma t}$ achieves identification, we can ask for an optimal Γ. The Γ we wrote down is actually optimal, as formalized below.

Proposition 4 (Optimal weights Γ for the GIV $y_{\Gamma t}$). *Given some GIV weights Γ_i with $\sum_i \Gamma_i = 0$, consider the GIV $z_t = y_{\Gamma t} = \sum_i \Gamma_i y_{it}$. Then, the asymptotic variance of the estimator α_T ($\sigma_\alpha^2 = \lim_{T\to\infty} Tvar(\alpha_T - \alpha)$) satisfies $\sigma_\alpha^2(\Gamma) = \frac{\sigma_\varepsilon^2}{\mathbb{E}[y_{St}y_{\Gamma t}]^2/\mathbb{E}[y_{\Gamma t}]}$. The value*

$$\tilde{\Gamma} = S - \tilde{E}, \qquad \tilde{E} \coloneqq \frac{V^{-1}\iota}{\iota' V \iota},$$

gives the optimal GIV estimator, in the sense that for any other Γ that is not collinear to $\tilde{\Gamma}$, the asymptotic variance $\sigma_\alpha^2(\Gamma)$ is larger. Then the shocks are iid, this implies $\tilde{E}_i = \frac{1}{N}$, and when they are uncorrelated, this implies $\tilde{E}_i \coloneqq \frac{1/\sigma_{u_i}^2}{\sum_j 1/\sigma_{u_j}^2}$

Hence, the "essence" of the GIV is not to be "size weighted minus value weighted idiosyncratic shocks", but rather "size weighted minus precision weighted (i.e. inverse-variance weighted when shocks are uncorrelated) idiosyncratic shocks.

3 Granular Identification: General Setup and Procedure

The previous section introduced the GIV in a simple context, with 0 or 1 loops. Here we generalize the framework to admit an arbitrary number of loops and factors.

[16] And we will have $\varepsilon_t^y = \eta_t + u_{Et}$.

3.1 Setup

We now propose a more general setup with potentially several factors, arbitrary loop structure and rich heterogeneity. Consider the following model of stationary "actions" y_{it} (e.g. employment, investment, TFP shock, return etc) by "entity" i (e.g. firm i in a closed-economy setting, or country i in an international setting):

$$y_{it} = u_{it} + \sum_f \lambda_{it}^f F_t^f + k_i + m' C_{it}, \tag{62}$$

where C_{it} is vector of controls that may include lagged demands and other characteristics, u_{it} are idiosyncratic shocks, and each F_t^f is a factor. It follows:

$$F_t^f = \eta_t^f + \alpha^f y_{S^f,t} + k^f + \phi^f C_t. \tag{63}$$

It depends on an exogenous shock η_t^f, and some on the mean action $y_{S^f,t}$, with weights that might depend on the factor; and on a set of controls C_t. Those controls could be e.g. lagged values, and include a 1.[17] We assume that all "size" factors have been normalized to add up to 1, $\sum_i S_i^f = 1$.

We partition the set of factors into observables factors $\mathcal{F}^O$ (e.g., GDP, aggregate investment, concentration) and hidden factors $\mathcal{F}^H$ (for some latent factors).

We also partition them into "exogenous factors" $\mathcal{F}^{\text{Exo}}$ where we know that $\alpha^f = 0$, and "endogenous" factors $\mathcal{F}^{\text{Endo}}$, where α^f may be non-zero α_f.

We call $\mathcal{F}^{\text{Exo},O}$ the set of observable, exogenous factors – those will serve as controls. The extended set of controls is

$$\mathcal{C}_t = \left(C_t, \left(F_t^f \right)_{f \in \mathcal{F}^{\text{Exo},O}} \right). \tag{64}$$

In some cases, a previously hidden factor can be made observable by some simple procedure — then we include it in the set of observable factors.[18]

We model the exposures to factors as

$$\lambda_{it}^f = \lambda_0^f + \lambda_1^f x_{it}^f + \zeta_i^f, \tag{65}$$

with $x_{it}^f \in \mathbb{R}$ is an observable (e.g. x_{it} is the de-meaned log size of entity i) and we normalize $x_{Et} = 0_k$ and $\zeta_E = 0$. We will normalize $\lambda_0^f = 1$, or $\lambda_1^f = 1$. Indeed, as as only $\lambda_{it} F_t$ matters, hidden factor f exhibit a 1-dimensional rescaling indeterminacy (i.e., the model doesn't change if we can multiply the factor by 10, and divide λ_{it} by $\frac{1}{10}$).

[17] They are observable and exogenous, so could fit in the general formalism as a factor – but it is clearer expositionally to separate them.

[18] This is the case in our initial simple example of Section 2.2, where $y_{Et} = \eta_t$ (up to a $O\left(\frac{1}{\sqrt{N}}\right)$ term) so that η_t is really an observable factor. Likewise, in the case of Section 4.1 with no feedback loop from sovereign yields to banks, r_{Et} is an exogenous, observable factor.

We call $\mathcal{F}^{\lambda_1}$ (resp. $\mathcal{F}^{\zeta}$) the set of factors for which we do not assume that $\lambda_1^f = 0$ (resp. $\zeta_i^f = 0$). When a factor is hidden, we will need a λ_1^f to be non-zero to make is visible and measure ζ_i^f. Hence, we assume if f is a hidden and $f \in \mathcal{F}^{\zeta}$, then $f \in \mathcal{F}^{\lambda_1}$.

We make the following identifying assumptions, for all f, i, the noise u_{it} are idiosyncratic:

$$u_{it} \perp \left(\eta_t^f, C_t, x_{it}^f\right), \tag{66}$$

but the η_t^f may be correlated across f's. The u_{it} may have some correlation (Section 19.9).

Multipliers Solving for the model gives, $y_{St} = \sum_f \lambda_S^f F_t^f + u_{St}$, i.e.

$$y_{St} = \sum_f \lambda_S^f \left(\eta_t^f + \alpha^f y_{St}\right) + u_{St}. \tag{67}$$

So, we can solve for the aggregate outcome y_{St} as: $y_{St} = \frac{u_{St} + \sum_f \lambda_S^f \eta_t^f}{1 - \sum_f \lambda_S^f \alpha^f}$, i.e.

$$y_{St} = M\left(u_{St} + \sum_f \lambda_S^f \eta_t^f\right), \tag{68}$$

where the multiplier

$$M = \frac{1}{1 - \sum_f \lambda_S^f \alpha^f} \tag{69}$$

measures the total impact of shocks, after feeding from the feedback loops:

$$M = \frac{dy_{St}}{du_{St}}.$$

Hence, an idiosyncratic shock has an impact that M times bigger than its direct effect.

Also, the total impact of an idiosyncratic shock on factor f is:

$$\frac{dF_t^f}{du_{St}} = M\alpha^f. \tag{70}$$

Our regressions will allow to identify M and $M\alpha^f$.[19]

[19] In addition, $\frac{dy_{St}}{d\eta_t^f} = M\lambda_S^f$ and $\frac{dF_t^g}{d\eta_t^f} = M\alpha^g\lambda_S^f + 1_{f=g}$. We also note the more advanced impact of a shock to the variable $j \neq i$:

$$\frac{dy_{it}}{du_{St,-i}} = M\sum_f \lambda_i^f \alpha^f, \tag{71}$$

which can also be estimated, see (174). The above equation is the reaction dy_{it} of to a shock du_{St}, when we also impose that $du_{it} = 0$.

3.2 An Identifiability Result

We rewrite model (62)-(63) in vector form:

$$y_t = \Lambda_t F_t + u_t + C_t m, \tag{72}$$
$$F_t = \alpha y_{St} + \eta_t + C_t \phi^f, \tag{73}$$

with Λ_t a $N \times r$ matrix, F_t a $r \times 1$ vector, C_t an $N \times c$ matrix (so C_{it} has c components, for typically c a small number say less than 5), m is $c \times 1$.[20]

The first column of Λ_t is assumed to be a vector of 1's, ι. In this section, we assume that we know the shape of the factor loadings, i.e. know the Λ_t, up to a rotation $\Lambda_t R$ (for R an invertible $r \times r$ matrix). For instance, depending on the application we know that the loading is an affine function of log GDP/Capita, or log Market Cap, or the stock market beta of a bank. This is consistent with the practice in modern finance (see e.g. Fama and French (1993)). Later we develop the full non-parametric version, estimating the factors. We don't have a priori information about the η_t, nor their covariance V^η.

To obtain identification, we need that Γ is not 0, and somehow that V^u can be recovered. In what follows, we give conditions to ensure this. We shall make the following two assumptions.

Assumption 1 (Condition for identification with GIV) *The vector $V^u S$ is not spanned by the factors loadings λ^f (where V^u is the covariance matrix of u_t).*

Assumption 2 (Restriction on the admissible variance-covariance matrix of residual u_t) *The variance-covariance on u_t is diagonal.*

Assumption 1 is essential and could not be relaxed. It ensures that the GIV is not identically 0. Economically, this assumption seems like a mild restriction. It is generically satisfied.[21]

Assumption 2 could be relaxed in number of ways.[22] Other sufficient condition for identification might be that V^u is $k-$sparse, e.g. has at most k non-zero off-diagonal elements, for some k, e.g. $N - r^2$ (see also Zou et al. (2006)). Another is to allow for some correlation that depends on the distance between entities i and j, perhaps via Gaussian processes (Rasmussen and Williams (2005)). We leave for extension to future research.

We next state a formal identification result.

[20] Our initial examples are particular cases of the general procedure, as detailed in Section 8.3.

[21] One thing that is does prevent is the case where the variance would be inversely proportional to size: then, GIV would fail, as then $V^u S = a\iota$ for some scalar a. Fortunately, in most contexts, variance may decay a bit with size, but less violently than one over size (see e.g. Lee et al. (1998) and the discussion in Gabaix (2011)).

[22] However, relaxations of Assumption 2 will still need to ensure some restriction on the space of variance-covariances allowed.

Proposition 5 (Sufficient condition for identification with GIV) *Consider the factor model above, when N is fixed but $T \to \infty$, and makes Assumptions 1 and 2. Then, the model identifies M, α^f by GIV.*

Proof sketch The proof is in Section 8.1. Here we provide a proof sketch. For simplicity, we omit the controls C_t. We use a projection matrix Q (defined in (94)) satisfying:

$$Q\Lambda = 0. \tag{74}$$

and

$$QV^u \left(I - Q'\right) = 0. \tag{75}$$

Then, premultiplying (72) by Q, we have $Qy_t = Qu_t$. Next, we define $\Gamma \coloneqq Q'S$, and the GIV as the scalar:

$$z_t \coloneqq \Gamma' y_t, \tag{76}$$

i.e. $z_t = S'Qy_t = S'Qu_t = \Gamma' u_t$, i.e.

$$z_t = u_{\Gamma t}. \tag{77}$$

Assumption 1 ensures $\Gamma \neq 0$.

Recall that we have (68)

$$y_{St} = Mu_{St} + \varepsilon_t$$

for ε_t a shock correlated with the η_t but not with the u_{it}'s. Finally, we have

$$u_{St} = S'u_t = S'Qu_t + S'\left(I - Q\right)u_t = z_t + v_t,$$

with $v_t = S'\left(I - Q\right)u_t$. Now, (75) ensures $\mathbb{E}\left[z_t v_t\right] = 0$. Then, we can write:

$$y_{St} = Mz_t + \left(Mv_t + \varepsilon_t\right),$$

with $\left(Mv_t + \varepsilon_t\right)$ orthogonal z_t. Hence, we can estimate the multiplier M by OLS.

Likewise, we have (via (68) and (70))

$$F_t^f = \alpha^f Mz_t + \varepsilon_t^f$$

for some shock ε_t^f orthogonal to z_t. Hence, we can estimate $\alpha^f M$ by regressing F_t^f on z_t.

In both regressions, we can add the estimated common shocks η_t as controls, which improves the precision.

3.3 A step-by-step user's guide

We outline the benchmark procedure, alongside several extensions that may be useful in practice, for the following model[23]

$$y_{it} = u_{it} + \lambda_0 F_t + \lambda_{1i}\eta_t + \varepsilon_t + k_i + m'C_{it},$$

where the factor (e.g., a price) is:

$$F_t = \alpha y_{St} + \eta_t^F$$

and the loading λ_{1i} are

$$\lambda_{1i} = \lambda_1 + \lambda_2 x_i + \zeta_i$$

so that they can be parametric, non-parametric or semi-parametric.

When the idiosyncratic shocks are homoscedastic, we estimate $M = \frac{1}{1-\alpha\lambda_0}$ and α (and hence λ_0) as follows.

1. Estimate the panel regression

$$y_{it} = a_t + m'C_{it} + k_i + \delta_{it},$$

   where $\delta_{it}^e = y_{it} - a_t^e - m^{e'}C_{it} + k_i^e$ gives us an estimate of $u_{it} + \check{\lambda}_i'\eta_t$.

2. Estimate the factor, η_t, which can be done in two ways

   (a) Run period-by-period cross-sectional regressions of δ_{it}^e on the demeaned characteristic, $\check{x}_i = x_i - x_E$,

$$\delta_{it}^e = \check{x}_i\eta_t^x + u_{it}^x,$$

   where $\eta_t^{x,e}$ is an estimate of η_t.

   (b) Use a factor model or principal components analysis

$$\delta_{it}^e = \phi_i\eta_t^{PCA} + u_{it}^{PCA},$$

   where $\eta_t^{PCA,e}$ is an estimate of η_t.

3. Estimate the time-series regression

$$\delta_{\Gamma t}^e = \lambda_x^\Gamma\eta_t^{x,e} + \lambda_{PCA}^\Gamma\eta_t^{PCA,e} + z_t,$$

[23] The procedure in this section extends to the model with multiple endogenous factors, when the factor exposures, λ_i, depend on a vector of characteristics, and to the case where the characteristics may vary over time, x_{it}.

where z_t^e, the residual of this regression, is an estimate of $u_{\Gamma t}$. Depending on the cross-sectional dispersion in x_i compared to ζ_i,$\eta_t^{x,e}$ or $\eta_t^{PCA,e}$ is a more accurate estimate of η_t. We add both estimates as controls in constructing an estimate of $u_{\Gamma t}$. In some applications, economic theory guides the choice of x_{it} and $\zeta_i = 0$, and we skip step 2(b) and only use $\eta_t^{x,e}$. If no obvious choice of x_i is available, that is $\lambda_2 = 0$, then we skip step 2(a) and use $\eta_t^{PCA,e}$ only.

4. Estimate the multiplier, $M = \frac{1}{1-\alpha\lambda_0}$, by the regression

$$y_{St} = k + Mz_t^e + \beta^{S\prime}\mathcal{C}_t + e_t^S.$$

   where $\mathcal{C}_t = \left(C_{St}, \eta_t^{x,e}, \eta_t^{PCA,e}\right)$ is a vector of controls, which includes our estimated η_t. We refer to the estimated multiplier as M^e.

5. Estimate αM via the regression:

$$F_t = \alpha M z_t^e + \beta^{F\prime}\mathcal{C}_t + e_t^F.$$

   We refer to the estimated impact as $(\alpha M)^e$.

6. Estimate α as the ratio $\alpha^e \coloneqq \frac{(\alpha M)^e}{M^e}$. It is in estimating α that M needs to be estimated with sufficient precision. This is analogous to the weak instruments problem in the context of IV estimators.

When the idiosyncratic shocks are heteroskedastic, the estimation involves a loop. We start from $E_i = 1/N$. After step 3, we use the residuals to obtain an estimate of $\sigma_{u_i}^2$ and update the weights to $\tilde{E}_i = \sigma_{u_i}^{-2} / \sum_{i=1}^N \sigma_{u_i}^{-2}$. If $\zeta_i = 0$, then we use u_{it}^x, which we directly get from step 2. If we use both $\eta_t^{x,e}$ and $\eta_t^{PCA,e}$, we estimate

$$\delta_{it} = \phi_{i0}\eta_t^{x,e} + \phi_{i1}\eta_t^{PCA,e} + u_{it},$$

and use the residuals from this regression to estimate $\sigma_{u_i}^2$.

In all cases, our estimators are consistent for large N and T. However, to obtain consistency with finite N and large T, we need the precision-weighted $\tilde{E}$. This implies standard bias-efficiency trade-off if estimating volatilities reduces the efficiency of the estimator. If the idiosyncratic volatilities are related to size, we can parameterize them and estimate

$$\ln \sigma_{u_i}^2 = \sigma_0 + \sigma_1 \ln S_i + \epsilon_i,$$

and use $\sigma_{u_i}^2 = \exp(\sigma_0) S_i^{\sigma_1}$.

3.4 Discussion: Robustness to misspecification

The GIV procedure is robust to some forms to misspecification, and more fragile to others.

If we misspecify the vector S of sizes (e.g. define z_t using a wrong vector S), things are fine: z_t is still valid, it is just not optimal any more.

Likewise, if we misspecify the variance of the u_{it} (but keeping them uncorrelated), then things are essentially find: as $u_E = O\left(\frac{1}{N^{1/2}}\right)$, we need not have $\mathbb{E}\left[u_{\Gamma t} u_{Et}\right] = 0$, but the term $\mathbb{E}\left[u_{\Gamma t} u_{Et}\right]$ will still be small ($O\left(\frac{1}{N^{1/2}}\right)$), and negligible for large N.

The threat is that we might not control properly for common factors. In the limit, suppose that we do not do a factor analysis. Then, $z_t = y_{\Gamma t} = \lambda_\Gamma \eta_t + u_{\Gamma t}$, and we will fail to control for the λ_Γ factor. This is a small sample problem (with a large enough T we typically measure λ accurately).

4 Applications

4.1 Doom loops: Spillovers from banks to sovereign yields and vice-versa

4.1.1 Contagion from banks to sovereigns

We start from a model in which instability in the banking sector may affect the sovereign but not vice versa. While the model is misspecified if weakness of the sovereign spills over to the banking sector, the estimator is straightforward to implement and provides as a useful point of reference for the full feedback model. Section 9.1 provides a model microfoundation.

We model the return of bank i, r_{it}, as

$$r_{it} = \kappa_i + \beta_i F_t + u_{it},$$

where F_t is a common factor that impacts banking stocks with a heterogeneous exposure β_i.[24] We allow the financial sector to impact the sovereign bond yield via

$$\Delta y_t = c + \alpha_y r_{St} + \eta_t^y. \tag{78}$$

The model implies for value- and equally-weighted returns

$$r_{St} = \kappa_S + \beta_S F_t + u_{St},$$
$$r_{Et} = \kappa_E + \beta_E F_t + u_{Et},$$

and thus the residual from regressing value-weighted on equally-weighted returns uncovers $u_{\Gamma t}$. We write, using $u_{\Gamma t} \simeq u_{St}$,

$$\Delta y_t = c + \alpha_y u_{\Gamma t} + \zeta_t,$$

[24] F_t is a banking factor and hence results in u_{it} shocks that are more bank-specific than when using, for instance, the return on the aggregate Italian stock market.

which allows us to estimate α_y using OLS.

4.1.2 Financial doom loops

Next, we allow shocks to the sovereign also to affect the banking sector. To this end, we adjust the model to

$$r_{it} = \kappa_i + \left(\lambda_0^F + \lambda_1^F x_i\right) F_t + \left(\lambda_0^y + \lambda_1^y x_i\right) \Delta y_t + u_{it}, \tag{79}$$

where x_i is the (cross-sectionally de-meaned) log market cap.[25] F_t is again a general factor that is specific to the banking sector. The feedback from banks to sovereigns, equation (78), remains the same.

The model implies for value-weighted returns

$$r_{St} = \kappa_S + \lambda_S^F F_t + \lambda_S^y \left(c + \alpha_y r_{St} + \eta_t^y\right) + u_{St},$$

where we define $\lambda_S^F \coloneqq \lambda_0^F + \lambda_1^F x_S$, $\lambda_S^y \coloneqq \lambda_0^y + \lambda_1^y x_S$ and hence

$$r_{St} = M\left(\kappa_S + c\lambda_S^y\right) + M\lambda_S^F F_t + M\lambda_S^y \eta_t^y + M u_{St},$$

where

$$M = \frac{1}{1 - \lambda_S^y \alpha_y}.$$

This multiplier M is a measure of the "doom loop" (Farhi and Tirole (2017)). It indicates how a "primitive" shock of 1% to the bank returns affects the sovereign yield (with a strength α_y), which in turns affect the banking sector (with a strength λ_S^y) with a strength M%.

For yield changes, we have

$$\Delta y_{St} = c + \alpha_y M u_{St} + \xi_t.$$

So that the total impact of banking shock on the sovereign yield is is $\alpha_y M$, i.e. the direct impact α_y times the doom loop M.[26]

To estimate the model, we proceed as follows:

1. Estimate a panel regression

$$r_{it} = \kappa_i + a_t + x_{i,t-1} b_t + v_{it},$$

where a_t and b_t are time fixed effects. We weigh the observation by the inverse of the variance estimated using a 3-month rolling window. We use v_{it} to form v_{St}. [27]

[25] We obtain similar results when we use a bank's volatility as conditioning variable.

[26] We could imagine other channels: for instance, bank return that lost money less less, which lowers GDP, weakens government finances and increases the yield. Our multiplier M and α_y should be understood as inclusive of all those indirect channels. [1/17/19 notes for an analytical example].

[27] The results are very similar of we use two instead of three months. However, given the large changes in volatility during this period it is important to allow for time variation in the volatility estimates.

2. We estimate the multiplier, M, using the OLS regression

$$r_{St} = c + Mu_{St} + \zeta_t^r.$$

3. We estimate $\beta = \alpha_y M$ using the OLS regression

$$\Delta y_{St} = c + \beta u_{St} + \zeta_t^y.$$

We then have $\lambda_S^y = (M-1)/\beta$ and $\alpha_y = \beta/M$.

4.1.3 Data

We estimate the model using data from the Italian banking sector and government bond market.[28] We construct a list of banks using ORBIS from Bureau van Dijk. We keep the banks that have been listed (but may delist during the sample). We merge the banks using ISIN codes to the research lists for equities from Thomson Reuter's Datastream[29] and only keep the primary security. We drop banks with a market capitalization smaller than 100 million euros.[30] We plot the number of banks at any given point in Figure 1. The number of banks declines during the European sovereign debt crisis and then stabilizes. In all periods, our sample includes more than 20 banks. We do note, however, that some banks are part of financial conglomerates (for instance, Banca Generali S.p.A. is an Italian bank that is part of the Generali Group).

In Figure 2, we plot the concentration using the banks' market equity, highlighting the high concentration in the sector at all times.

We use sovereign 10-year benchmark yields from Thomson Reuters. We compute arithmetic returns for stocks and simple differences in yields as our measure of Δy_t. We use weekly returns at a daily frequency (using overlapping data) to mitigate the impact of non-synchronous trading. However, using daily returns results in similar, and in fact somewhat larger, estimates.

In Figure 3, we plot the dynamics of German and Italian yields since the introduction of the Euro. During the first period, Italian yields and Germany yields move in lockstep following the introduction of the Euro. This changed drastically following the financial crisis and in particular the European sovereign debt crisis when yields diverged.

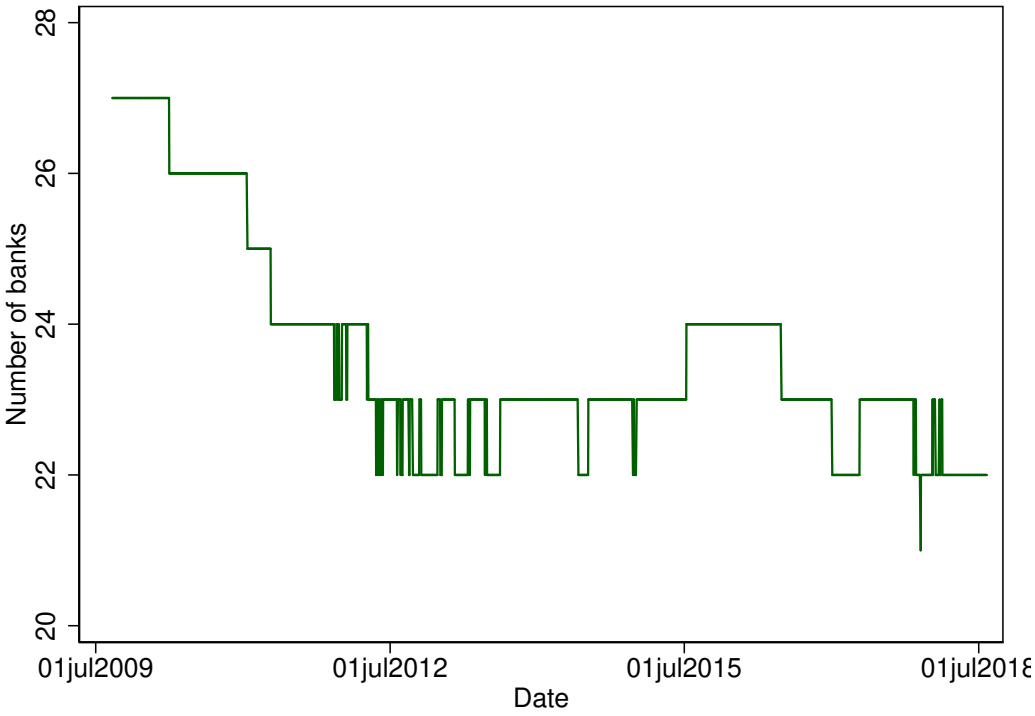
In Figure 4, we plot the dynamics of yields and value-weighted returns. The figure clearly illustrates the strong negative correlation between the sovereign yield and the returns on the banking sector during the European sovereign debt crisis. The same happened in recent months following the political uncertainty in Italy in May 2018.

[28] We are in the process of extending this application to other countries in the Euro area.

[29] For Italy, the codes for these lists are FITA and DEADIT, and include both alive and dead securities.

[30] We also remove a small fraction of observations with daily returns over 100% or for which the total return index falls below one to avoid problems with discreteness. Despite these screens, a concern that remains is that some of the banks in our sample are small, which may lead to liquidity problems.

Figure 1: Number of banks. The figure plots the number of banks listed in Italy with a market capitalization larger than 100 million euros. We use data from September 2009 to July 2018.



4.1.4 Empirical results

The estimation results are reported in Table 1. In Panel A, we report the results for the sample from September 2009 until July 2018, while Panel B reports the results from December 2014 until July 2018. The European Central Bank announced the large-scale asset purchase programme in December 2014, which may have helped to stabilize financial markets and weaken the doom loop.

In the first column, we report the results of an OLS regression of (weekly) yield changes on value-weighted banking returns. The coefficient is -0.015 in the full period and -0.006 in the second, and both are statistically significant. The R-squared during the full sample is 19%, highlighting substantial comovement. However, the OLS estimate likely reflects common factors that affect both the sovereign and the banking sector, and this correlation tells us little about financial doom loops.

In the next two columns, we estimate the impact of weakness in the banking sector on sovereign yields ("GIV partial feedback"). We first regress value-weighted returns on equally-weighted returns to obtain the idiosyncratic component of large banks' returns. The R-squared is around 90%, indicating that 10% of the large-cap banking returns are idiosyncratic. The coefficient is 1.29 (full sample) and 1.17 (second sample), indicating that large banks have a higher beta than small banks. This is consistent with evidence from the U.S. banking sector, see Gandhi and Lustig (2015). A reason is that large banks have higher leverage compared to small banks, which is the case in our sample as well.

Figure 2: Concentration in the banking sector. The figure plots the square root of the Herfindahl index from September 2009 until July 2018. The Herfindahl index is calculated using banks' market equity.

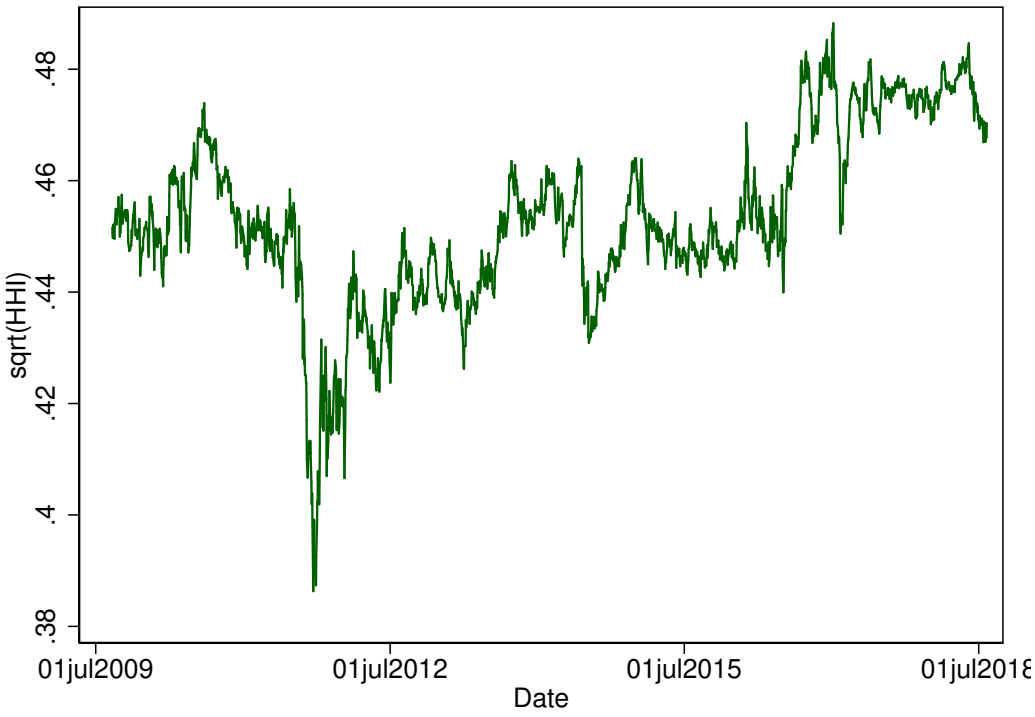


Table 1: Linkages between banks and sovereigns. The table reports the estimates of the link between the banking sector and sovereign yields in Italy during the European sovereign debt crisis and its aftermath. We estimate the models in Panel A using data from September 2009 to July 2018 and in Panel B using data from December 2014 to July 2018. The first model is an OLS regression from yield changes on value-weighted banking returns. The second model estimates the causal impact using GIVs from banks to the sovereign. The third model estimates the full feedback loop using GIVs.

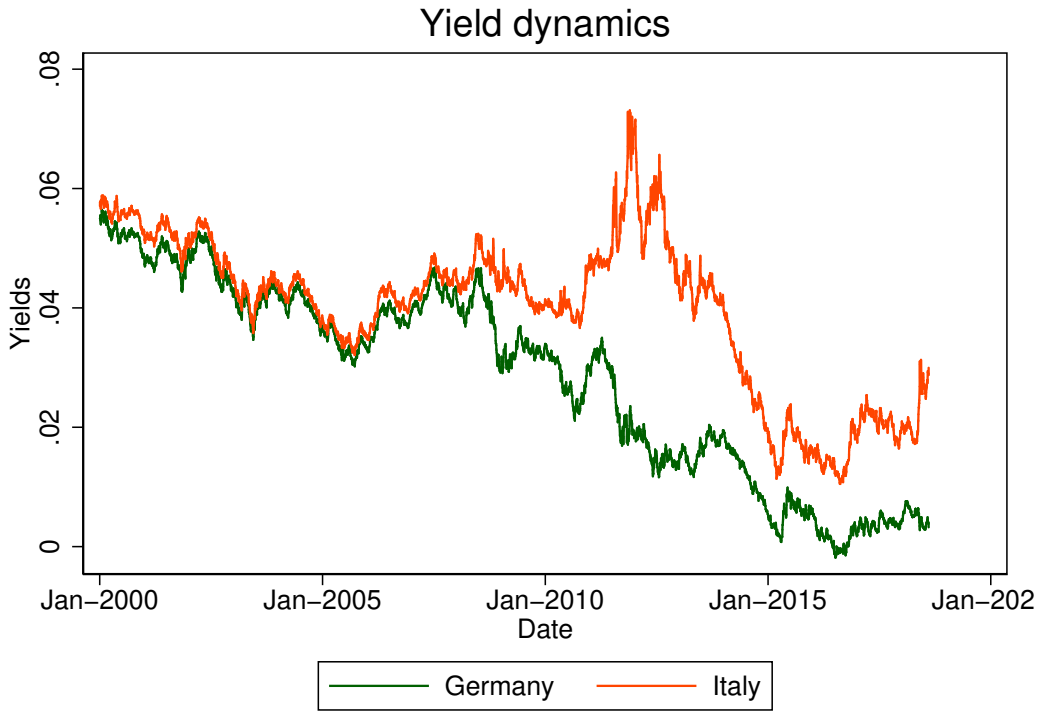
Panel A: 2009.9 - 2018.7

	OLS	GIV partial feedback		GIV full feedback	
			Banks→ Sovereign	Sovereign → Banks	Banks → Sovereign
	Δy	r_S	Δy	r_S	Δy
r_S	-0.015				
t-stat	(-9.56)				
r_E		1.29			
t-stat		(55.13)			
u_Γ			-0.011	1.45	-0.015
t-stat			(-2.67)	(11.95)	(-2.93)
R^2	19.3%	88.9%	1.2%	16.60%	1.60%

Panel B: 2014.12 - 2018.7

	OLS	GIV partial feedback		GIV full feedback	
			Banks → Sovereign	Sovereign → Banks	Banks → Sovereign
	Δy	r_S	Δy	r_S	Δy
r_S	-0.006				
t-stat	(-3.59)				
r_E		1.17			
t-stat		(35.77)			
u_Γ			-0.0020	1.15	0.0012
t-stat			(-0.59)	(4.33)	(0.30)
R^2	5.3%	89.3%	0.7%	5.9%	0.1%

Figure 3: Yield dynamics in Germany and Italy. The figure plots the 10-year government bond yields from January 2000 until July 2018 in Germany and Italy.

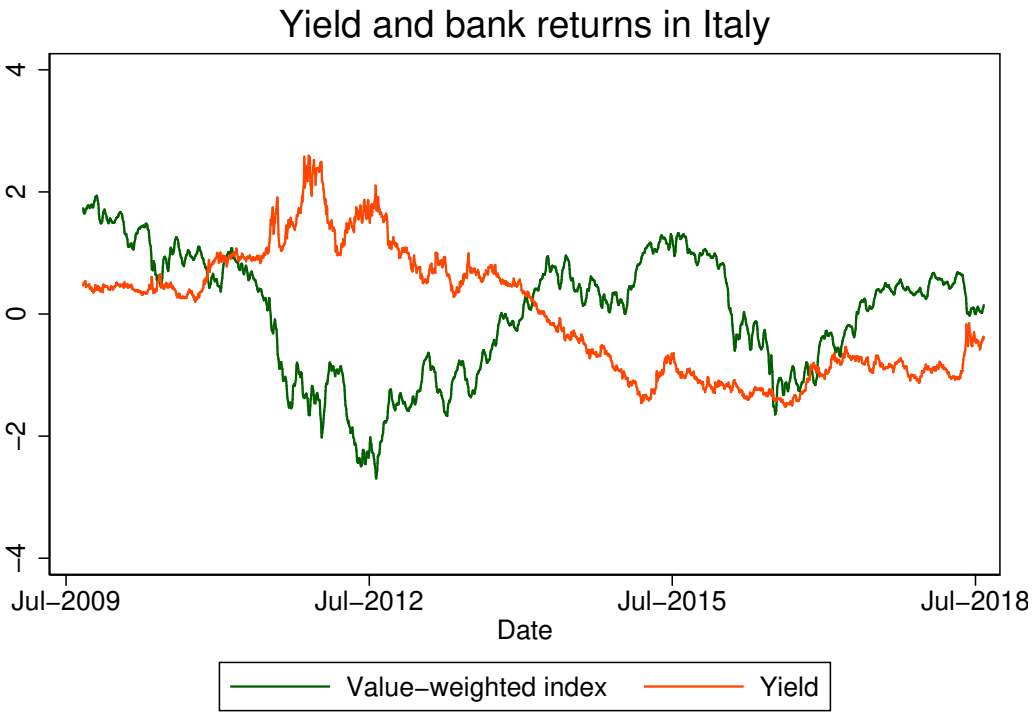


In the next column, we report the causal impact of banks on sovereigns, and estimate a coefficient for α_y of -0.011 ($t = -2.67$), which is somewhat smaller than the co-movement implied by the OLS estimate. In the second sample, after the sovereign debt crisis, the GIV estimate is substantially smaller than the OLS estimate and insignificant. Hence, we find no evidence that banks impact sovereign yields during this calmer period.

In the last two columns, we estimate the full feedback loop. During the full sample, we estimate a significant feedback M of 1.45, which remains significant but smaller at 1.15 during the second sample. So, primitive shocks are amplified with a multiplier of M. To the best of our knowledge, this is the first causally estimated measure of doom loops.

The impact of banks on the sovereign during the full sample equals -0.015, but as explained before, this coefficient corresponds to β. To uncover α_y, we use $\alpha_y = \beta/M = 0.010$, which is close to the GIV estimate coming from the partial feedback model (but lower than the OLS estimate). This coefficient $\beta = -0.015$ means that a 50% decrease in bank stock returns adversely affect the sovereign yield by increasing it by 75 basis points. During the second sample, the estimate is insignificant again. These estimates together imply point to a financial doom loop that was operative during the European sovereign debt crisis. Going forward, one could imagine drawing rolling window measures of doom loops.

Figure 4: Sovereign yield and banking returns. The figure plots the 10-year government bond yield from September 2009 until July 2018 in Italy alongside the equity return on the value-weighted banking index.



4.2 Oil supply elasticity

4.2.1 Data

Our data construction follows the existing literature (Kilian 2009; Caldara et al 2018). The data on oil supply and prices are from the U.S. Energy Information Administration (EIA). We observe the monthly oil supply for 21 countries (OPEC and non-OPEC) from January 1985 until December 2017.[31] The real oil price series is obtained based on the refiner acquisition cost of imported crude oil and deflated using the US CPI to obtain the real price of oil. We focus on short-run supply elasticities and use monthly changes in supply and prices.

Discuss use of truncation

4.2.2 Model

The oil supply curve is given by

$$y_{it} = k_i + \lambda_0 p_t + \lambda_{it}\eta_t + u_{it},$$

where $\lambda_{it} = \lambda_0 + \lambda_i x_{it} + \zeta_i$, where $x_i = 1$ is a country is a member of OPEC.[32] Price in turn depends on aggregate demand and supply,

$$p_t = \alpha y_{St} + \epsilon_t.$$

4.2.3 Estimates

...

4.3 Aggregate GDP and employment multipliers

4.3.1 Model

Our specification and data construction follow Nakamura and Steinsson (2014), but we extend the model to allow for aggregate spillovers. Let y_{it} denote 2-year growth rate in per-capita employment or real per-capita GDP in state i from year $t-2$ to year t.[33] The model is

$$y_{it} = k_i + \alpha y_{St} + \lambda_i \eta_t + u_{it}, \tag{80}$$

where $\lambda_i = \lambda_0 + \lambda_1 x_i + \zeta_i$. We define the multiplier as $M = (1-\alpha)^{-1}$. As x_i measures the exposure to aggregate shocks, we use the average growth rate in state i during the four worst declines in aggregate growth, y_{St}, in our sample.

[31] We follow Cadara et al. (2018) and remove Gabon from the sample due to concerns about data quality.

[32] As OPEC membership changes over time, x_{it} is indexed by time.

[33] Alternatively, we could model the annual growth rate and add lags to capture the persistence in GDP. To make our estimates comparable to the existing literature we use 2-year growth rates instead.

Table 2: State-level growth spillovers. The table reports the multiplier estimates for real per-capita output growth and per-capita employment growth using data from 1963 to 2017. We use the model in (80) to estimate the multiplier.

Output	M	$\eta_t^{x,e}$	$\eta_t^{PCA,e}$	R-squared
Estimate	2.52	-0.11	0.18	0.62
St. error	0.45	0.01	0.05	
Employment	M	$\eta_t^{x,e}$	$\eta_t^{PCA,e}$	R-squared
Estimate	1.63	-0.05	-0.25	0.59
St. error	0.49	0.01	0.03	

4.3.2 Data

All growth rates are computed on a per-capita basis. For GDP, we start from annual nominal GDP growth using data from the Bureau of Economic Analysis. We deflate the nominal GDP numbers using the national CPI. We construct per-capita growth rates by dividing the real GDP numbers by the state-level population numbers from the U.S. Bureau of the Census. We use nominal GDP or employment as weights to aggregate growth rates. We use data from 1963 to 2017, providing us with growth rates from 1965 to 2017. We remove the average growth rate and a time trend from each of the series. In addition, we remove four states (Alaska, North Dakota, South Dakota, and Wyoming) with extreme growth rates in the early years of the sample, see also Sarto (2018).

4.3.3 Estimates

We report the point estimates in Table 2 for output and employment where we use $\eta_t^{x,e}$ and $\eta_t^{PCA,e}$ as controls.[34] The point estimate for GDP growth is somewhat larger at 2.5 than for employment growth at 1.6, but the confidence intervals overlap. To be conservative, we use $M = 1.5$ in our calculations below.

4.4 Industry concentration and corporate investment

4.4.1 Data

We use data constructed by Gutiérrez and Philippon (2017b) (GP), who start from Compustat data and make a series of adjustments as Compustat only contains a subset of all U.S. firms and does not account for imports. We focus on $CR_{jt}(n)$ as our concentration measure, which is the share of sales by the largest n firms in industry j. While this measure measure is perhaps ad-hoc in that

[34] Including the factors as controls is quite important in this application, instead of simply regressing y_{St} on $y_{\Gamma t}$. If we do not account for the factors, the multiplier estimates are much smaller, which indicates that smaller states are more exposed to aggregate shocks. In addition, adding the factors as controls sharpens the estimates.

it selects a discrete cutoff instead of a more smooth measure that accounts for all firms, it requires fewer additional assumptions than, for instance, the Herfindahl index as we do not have firm-level data for the entire population of firms.

Starting from the sales share of the top n firms in an industry based on Compustat data, GP multiply it with the share of Compustat sales relative to the sum of gross output and imports for a given industry. If import data are unavailable, as is the case for services, the Compustat share is adjusted using gross output data only. The net investment rate, ni_{jt}, is defined as gross investment minus depreciation relative to capital in the previous period. The gross and net investment rate are constructed using data from the U.S. Bureau of Economic Analysis (BEA). The BEA data are from 1995 to 2015. We follow the same industry classification using BEA segments as GP, and we only include industries for which the number of firms exceeds five in all years. This results in 41 industries in total.

4.4.2 Facts and OLS results

As documented by GP and others, concentration has increased during the last two decades while investment declined. In Figure 5, we plot the dynamics of $CR(8)$ alongside the net and gross investment rates. We aggregate all measures across industries using gross output.

In addition to the aggregate trends, variation in concentration across industries also appears to be related to changes in investment rates. In Figure 6, we plot the concentration measures in 1995 and 2015 as well as the net investment rates. A cross-sectional regression of the low-frequency change in the net investment rate on the change in concentration results, while controlling for the change in Tobin's Q, in a regression coefficient of -0.10 ($t = -2.78$). [35]

We estimate a panel regression to relate concentration and investment rates

$$ni_{jt} = \kappa_j + \beta CR_{jt}(n) + \delta Q_{j,t-1} + \epsilon_{jt}, \tag{81}$$

where we control for the lagged Tobin's Q, add industry fixed effects, and add year fixed effects in some specifications, implying that we identify β from variation within industries only. We weight industries by the average size of the capital stock.

In Table 3, we report the results from the OLS regression. This OLS regression does not instrument changes in industry concentration, and it is not a causal analysis, unlike the GIVs that we will next show. In the first two columns, we only include industry fixed effects, but no year fixed effects. In the first column, we do not cluster the standard errors, while in the second column they are clustered by industry. In the third column, we include year fixed effects and cluster standard errors by industry. The results, if they were taken at face value (which one should not, as they are not instrumented) might suggest that a 1% increase in the concentration ratio results in a 0.1-0.07%

[35] If we weight the observations by the size of the average capital stock of an industry or by the absolute change in concentration, the slope coefficients change to -0.10 ($t = -3.77$) and -0.092 ($t = -4.19$), respectively.

Figure 5: Trends in concentration and investment rates. The figure displays aggregate trends in concentration as measured by the $CR(8)$ measure (top left panel), the net investment rate (top right panel), and the gross investment rate (bottom left panel) from 1995 to 2015. Industry-level data are aggregated using gross output.

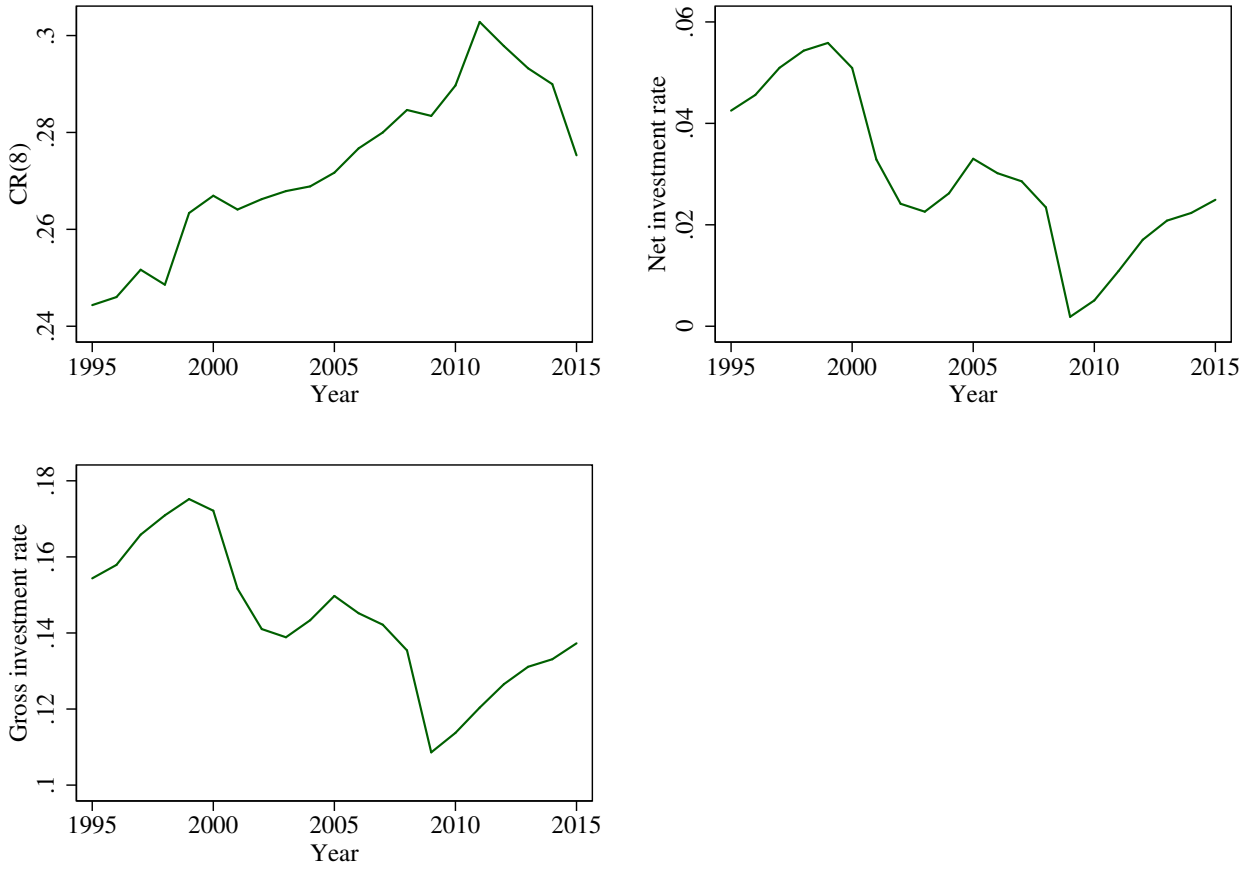
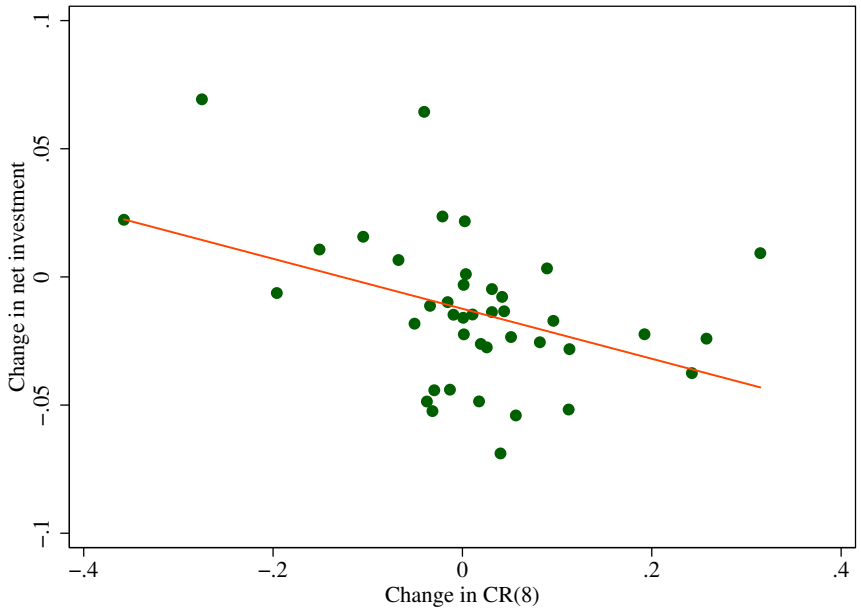


Figure 6: Changes in concentration and investment across industries. The figure displays the change in concentration as measured by the $CR(8)$ measure from 1995 to 2015 (horizontal axis) and the change in the net investment rate over the same period (vertical axis) across industries.



decline in the net investment rate.

We vary the concentration measure to $C_{jt}(n)$, $n \in \{4, 8, 20, 50\}$, in Table 4. The regressions include industry and year fixed effects, and we cluster standard errors by industry. The results are strongest for the $C(8)$ and $C(20)$ measures, but the results are broadly consistent across specifications.

4.4.3 GIV results

Next, we use GIV to provide a causal estimate of β under different assumptions about the drivers of firm-level growth rates.

Modeling growth rates and constructing an instrument The concentration measure is defined by

$$C_{jt}(n) = \sum_{i=1}^{n} S_{ijt} = \sigma_{jt} \sum_{i=1}^{n} S_{ijt}^{CP} = \sigma_{jt} C_{jt}^{CP}(n),$$

where $n < N$, σ_{jt} is the share of Compustat sales relative to gross output and potentially imports, S_{ijt}^{CP} is the sales share based on Compustat data alone, and $C_{jt}^{CP}(n)$ is the concentration measure

Table 3: Net investment and concentration. The table shows the OLS estimates of (81) using data from 1995 to 2015. The specifications in the first two columns do not include year fixed effects while the specification in the third column does. Standard errors are clustered by industry in the second and third columns. We use the $CR(8)$ measure to measure concentration.

	I/K(t)	I/K(t)	I/K(t)
CR(8,t)	-0.100	-0.100	-0.070
	(-6.78)	(-2.96)	(-2.53)
Q(t-1)	0.019	0.019	0.013
	(9.48)	(3.46)	(2.91)
Time F.E.			✓
Observations	861	861	861
R^2	0.154	0.154	0.280

t statistics in parentheses

Table 4: Net investment and concentration: Robustness. The table shows the OLS estimates of (81) using data from 1995 to 2015, but we vary the concentration measure to $C_{jt}(n)$, $n \in \{4, 8, 20, 50\}$. All specifications include year fixed effects and standard errors are clustered by industry.

	I/K(t)	I/K(t)	I/K(t)	I/K(t)
Q(t-1)	0.013	0.013	0.013	0.013
	(2.95)	(2.91)	(2.94)	(2.98)
CR(4,t)	-0.054			
	(-1.75)			
CR(8,t)		-0.070		
		(-2.53)		
CR(20,t)			-0.067	
			(-2.60)	
CR(50,t)				-0.055
				(-2.01)
Time F.E.	✓	✓	✓	✓
Observations	861	861	861	861
R^2	0.270	0.280	0.280	0.273

t statistics in parentheses

using only Compustat data. We approximate the 5-year change in the concentration measure as

$$\Delta C_{jt}^{(5)}(n) \simeq \sigma_j \sum_{i=1}^{n} S_{ij,t-5}^{CP} \left(g_{ijt} - g_{Ejt} \right), \tag{82}$$

where g_{ijt} is the sales growth of firm i in industry j over a 5-year period[36] and $\sigma_j = \overline{\sigma_{jt}}$. This expression highlights that shocks that affect all firms in a given industry identically do not affect concentration. To mitigate the impact of outliers, we use geometric growth rates and winsorize them at the 1% and 99% level by industry. [37]

We consider three models of firm-level growth rates that may help to understand the channel via which concentration relates to net investment rates. In the first model, growth rates evolve as

$$g_{ijt} = \alpha_{jt} + \lambda x_{ij,t-5} \eta_t + u_{ijt}, \tag{83}$$

with the normalization $x_{Ejt} = 0$.

If x_{ijt} is demeaned log size, a positive $\lambda \eta_t$ corresponds to a shock (e.g. a regulatory shock) that favors the growth of large firms, and penalizes the growth of small firms — e.g. perhaps because large firms can adopt new technologies better, because of changes in anti-trust regulation, or more darkly because they can lobby.

The industry-year effect will be absorbed in computing $g_{ijt} - g_{Ejt}$. We compute the instrument as

$$z_{jt}(n) = \sigma_j \sum_{i=1}^{n} S_{ij,t-5}^{CP} u_{ijt}.$$

In this model, we isolate the variation in firm-level growth rates adjusting for an aggregate shock that affects all industries in the same way. Within industries, however, small and large firms may be affected differentially.

In the second model, we allow for common shocks that differ for manufacturing and non-manufacturing industries,

$$g_{ijt} = \alpha_{jt} + \lambda_{\text{Manuf}(j)} x_{ij,t-5} \eta_{\text{Manuf}(j)t} + \dot{u}_{ijt}, \tag{84}$$

which we estimate by regression firm-level growth rates on x_{ijt} interacted with manufacturing-year fixed effects, where Manuf(j) groups industries by manufacturing and non-manufacturing. We construct the instrument by replacing u_{ijt} by $\dot{u}_{ijt}$ in calculating the instrument.

In the third model, we allow for industry-specific shocks that differentially affect small and large

[36] We omit firms that enter and exit the sample during this period, which may weaken the instrument. An alternative approximation would use growth rates defined by $g_{ijt} = 2(Y_{ijt} - Y_{ij,t-5})/(Y_{ijt} + Y_{ij,t-5})$, as in Davis, Haltiwanger, and Miranda (2007), which is a way to account for entry and exit.

[37] To illustrate the impact of approximating σ_{jt} by σ_j, we run a panel regression of $\Delta C_{jt}^{(5)}(n)$ on $\sigma_j \Delta C_{jt}^{(5),CP}(n)$, including industry fixed effects. The R-squared of this regression is 51% with a slope coefficient of 0.96.

Table 5: Net investment and concentration. The table shows the OLS estimates of (86) using data from 1995 to 2015. The specifications in the first two columns do not include year fixed effects while the specification in the third column does. Standard errors are clustered by industry in the second and third columns. We use the $CR(8)$ measure to measure concentration.

	I/K(t)	I/K(t)	I/K(t)
CR(8,t)-CR(8,t-5)	-0.096	-0.096	-0.075
	(-5.90)	(-3.21)	(-2.72)
CR(8,t-5)	-0.11	-0.11	-0.060
	(-5.68)	(-2.18)	(-1.48)
Q(t-1)	0.018	0.018	0.013
	(9.37)	(3.41)	(2.96)
Time F.E.			✓
Observations	861	861	861
R^2	0.154	0.154	0.281

t statistics in parentheses

firms

$$g_{ijt} = \alpha_{jt} + \lambda_j x_{ij,t-5}\eta_{jt} + \ddot{u}_{ijt}. \tag{85}$$

We compute the instrument as before but with $\ddot{u}_{ijt}$ replacing u_{ijt}.

Lastly, our instrument isolates the idiosyncratic variation in 5-year *changes* in concentration ratios using the approximation in (82). We rewrite the regression equation in (81) as

$$ni_{jt} = \kappa_j + \beta\Delta_5 CR_{jt}(n) + \gamma CR_{j,t-5}(n) + \delta Q_{j,t-1} + \epsilon_{jt}, \tag{86}$$

and instrument $\Delta_5 CR_{jt}(n)$ using the GIV. In Table 5, we first repeat the OLS estimation of this model, analogous to Table 3. We find, as expected, that the estimates of β and γ are similar. In addition, the OLS estimates imply that β is statistically significant.

Estimation results We consider a specification with and without time fixed effects. We include industry fixed effects in all specifications and cluster standard errors by industry. The first-stage estimation results are presented in Table 6.

We find that the instrument is in all three models strongly related to within-industry variation in concentration. In the first two columns, we only control for common shocks that differentially impact small and large firms. In the next two columns, we allow for different shocks to manufacturing and non-manufacturing industries. In the last two columns, we allow for industry-specific shocks that impact large and small firms differently. The second column of each pair includes time fixed effects. In all cases, the instrument is not weak and the lowest F-statistic is 29.2, which corresponds to the final specification (Column 6).

Table 6: First-stage estimation results. The table shows the first-stage estimation results of (86) using data from 1995 to 2015 where we use GIVs to instrument changes in concentration as measured by the $CR(8)$ measure. The results in the first two columns correspond to the firm-level growth model in (83), the next two columns to the model in (84), and the last two columns to the model in (85). The second column of each model includes year fixed effects. Standard errors are clustered by industry.

	Model 1	Model 1	Model 2	Model 2	Model 3	Model 3
Q(t-1)	-0.0023	-0.0030	-0.0024	-0.0030	-0.0038	-0.0036
	(-0.36)	(-0.40)	(-0.38)	(-0.39)	(-0.49)	(-0.43)
C(8,t-5)	-0.35	-0.36	-0.36	-0.36	-0.40	-0.41
	(-2.84)	(-2.75)	(-2.87)	(-2.74)	(-2.95)	(-2.94)
z(n), model 1	1.09	1.11				
	(10.68)	(9.72)				
z(n), model 2			1.05	1.07		
			(10.63)	(9.26)		
z(n), model 3					0.81	0.80
					(7.02)	(5.64)
Observations	861	861	861	861	861	861
R^2	0.446	0.464	0.433	0.453	0.371	0.393

t statistics in parentheses

Table 7: Second-stage estimation results. The table shows the second-stage estimation results of (86) using data from 1995 to 2015 where we use GIVs to instrument changes in concentration as measured by the $CR(8)$ measure. The results in the first two columns correspond to the firm-level growth model in (83), the next two columns to the model in (84), and the last two columns to the model in (85). The second column of each model includes year fixed effects. Standard errors are clustered by industry.

	Model 1	Model 1	Model 2	Model 2	Model 3	Model 3
C(8,t)-C(8,t-5)	-0.16	-0.14	-0.16	-0.14	-0.11	-0.100
	(-3.97)	(-1.87)	(-4.17)	(-1.88)	(-1.70)	(-1.05)
Q(t-1)	0.018	0.013	0.018	0.013	0.018	0.013
	(3.42)	(2.94)	(3.43)	(2.94)	(3.42)	(3.03)
C(8,t-5)	-0.14	-0.093	-0.14	-0.093	-0.12	-0.073
	(-2.90)	(-1.81)	(-2.98)	(-1.83)	(-2.09)	(-1.25)
Observations	861	861	861	861	861	861
R^2	0.136	0.266	0.137	0.266	0.153	0.279

t statistics in parentheses

The slope coefficient varies between 0.15 and 0.2, where part of the attenuation is due to the data limitations described before. With more complete data, including domestic firms that are not covered in Compustat and foreign firms, yet following the same steps, one may be able to strengthen the instrument further. At the same time, our instrument complements the literature by focusing on domestic shocks as opposed to increased competition from China (Autor et al. (2013)).[38]

In Table 7, we report the second-stage results. For the first two models, we find that investment is significantly negatively related to concentration, consistent with the OLS results and GP. The point estimates are in fact somewhat more negative compared to when we estimate the model using OLS.[39]

Take the coefficient of net investment on GIV-instrumented change in the concentration ratio, which is -0.16. It means that a 1% increase in the concentration ratio results in a 0.16% decline in the net investment rate.

The first conclusion from this application is that GIV can be used to construct instruments for changes in concentration that are not weak, and can plausibly be strengthened further with more granular data. Second, the link between concentration and investment is robust to common

[38]GP instrument concentration by "noisy entry" in the nineties, which affects concentration in the early 2000s. As they explain, their instrument relies heavily on an exceptional period in the late nineties in which entry is not well explained by fundamentals. GIV may be applicable outside of this episode as well, provided that concentration is sufficiently strong and idiosyncratic shocks are sufficiently large. Moreover, we can control for different types of aggregate shocks, as explained below, which may be helpful in understanding the channel.

[39]However, once we allow for common shocks at the industry level that may affect small and large companies differentially, the point estimate lowers and turns insignificant. This is probably because this control soaks up too much of any variation.

shocks to all industries that may impact small and large firms differently, or common shocks to manufacturing and non-manufacturing industries (but potentially with a differential impact on small and large firms). The link weakens when we allow for industry-specific shocks that may affect firms of different size differently.[40]

5 Inferring Macro Elasticities from Micro Elasticities

An energetic and ingenious recent literature has estimated "regional impacts". For instance, Mian and Sufi (2014) estimate a negative elasticity ψ^ℓ of local non-tradable employment to local housing prices. Nakamura and Steinsson (2014) estimate a local impact ψ^r of local government spending on local employment and GDP.

A key open question is: how to go from the local impact ψ^ℓ to the global impact ψ. We shall show that, under fairly reasonable conditions, we have the following. Using M as our GDP multiplier (estimated in Section 4.3), the global impact is:

$$\psi = M\psi^\ell. \tag{87}$$

We next justify this. We first sketch the assumptions (which are detailed in Section 8.4). There are feature N regions. In each region, consumers have Cobb-Douglas preferences in non-tradables and tradables.

There are shocks, in particular shocks about expectations of future income, which work as shocks to consumption. Crucially, there is no special policy to counteract the shocks under study, e.g. housing shocks. Hence they have the same propagation as other shocks. This is why our general "GDP multiplier" is legitimate to capture general equilibrium effects.[41] If each shocks had a different macro-policy, we'd need changes to the multiplier.[42]

As we focus on the short run, prices are sticky in the short run. More behaviorally, they have a high marginal propensity to consume and are non-Ricardian (they forget about future taxes).

Proposition 6 (Going from micro to macro estimates: impact of housing on employment). *Call ψ^ℓ the elasticity of local employment to local housing prices, as in Mian and Sufi (2014). Then, the aggregate impact (i.e. the elasticity of total employment to an aggregate change in housing prices)*

[40] There is a potential bias: if Apple grows a lot then the concentration in the consumer electronics sector goes up (as we used earlier). But also, Apple invests a lot (there is an idiosyncratic increase in net investment by Apple). So, purely via idiosyncratic shocks, we except a positive correlation between increase in large firm sizes (hence, concentration) and net investment. Hence, our coefficient β (called b in this section) will be upward biased. Section 11.1 discusses that, and finds that this bias is actually very small.

[41] An important enrichment is to distinguish between demand (i.e. preference) and supply (productivity) shocks, especially empirically, and have a multiplier M for demand shocks and another M for supply shocks. We plan to do that in a next iteration of this draft.

[42] Then, we could still do a GIV to estimate the multiplier M for those shocks. For instance, in the context of Proposition 6, we'd run a GIV relying on the weighted sum of idiosyncratic shocks to housing prices.

is:

$$\psi = M\psi^{\ell}. \tag{88}$$

The proof is in the appendix. For instance, Mian and Sufi (2014) find that a drop of 10% in housing net worth created a temporary fall in non-tradable employment of 2.1%, so i.e. a local elasticity of $\psi^{\ell} \simeq 0.21$.[43] As we find a GDP multiplier $M \simeq 1.5$, the global elasticity is $\psi = M\psi^{d} \simeq 0.3$. Given the aggregate fall in housing prices during 2006-9 was about 10% (Mian and Sufi (2014)), this accounts for an aggregate employment drop of $\psi \times 10\% = 3\%$.

Proposition 7 (Going from micro to macro estimates: impact of government spending on GDP). *Call ψ^{ℓ} the dollar impact on local GDP of \$1 of local government spending, as in Nakamura and Steinsson (2014) or Chodorow-Reich (2017). Then, the aggregate impact (i.e., the aggregate dollar impact on GDP of an increase of \$1 of government spending) is:*

$$\psi = M\psi^{\ell}\frac{\alpha^{NT}}{\alpha^{G,NT}}, \tag{89}$$

where α^{NT} is the share of non-tradables in the consumer's basket, and $\alpha^{G,NT}$ is the share of non-tradable in the government's basket.

Likewise, Nakamura and Steinsson (2014) find that \$1 of local government spending increases local GDP by $\psi^{\ell} \simeq \$1.4$. We assume that government and households spend equal shares on tradable vs non-tradable goods, $\frac{\alpha^{NT}}{\alpha^{G,NT}} = 1$. Given our estimate $M \simeq 1.5$, this leads to an aggregate elasticity of $\psi = M\psi^{\ell} \simeq 2$: a \$1 increase in government spending leads to an increase in aggregate GDP by \$2. This is a rather high estimate. It is consistent, however, with a marginal propensity to consume of $\mu = 0.5$, and non-Ricardian consumers, which would give $M = \frac{1}{1-\mu} = 2$ (see Gabaix (2018) and Section 8.4 for such a model with modern microfoundations).

6 Extensions of the framework

There are many ways to increase the setups in which the GIV idea applies.

One can handle multidimensional "actions": for instance, a firm could have a shocks that affect both productivity and labor demand. Formally, the actions y_{it} are now multidimensional. The GIV idea goes through, this is developed in Section 10.1.

The GIV can also be extended to non-homogeneous influences Suppose a model:

$$y_{it} = \gamma \sum_{j} G_{ij} y_{jt} + \lambda_i \eta_t + u_{it}, \tag{90}$$

[43] We calculated the 0.21 elasticity from their original data, and this is confirmed by the following back-of-the-envelope check. Analyzing the 2006-2009 period, they find that a local drop of 10% in housing created a temporary fall in non-tradable employment of 3.7%. Using a 0.7 share of nontradables (Nakamura and Steinsson (2014)), this would give an local elasticity of total employment to housing net worth of $\psi^{\ell} \simeq 0.37 \times 0.7 = 0.26$.

i.e., in vector form

$$y_t = \gamma G y_t + \Lambda \eta_t + u_t, \tag{91}$$

with a given "influence" matrix G (in our baseline model, $G = \iota S'$). We'd like to identify γ, the strength of linkages. The generalization of our GIV (which turns out to be optimal) is to define a "size" vector $S \coloneqq G'E$. Then, left-multiplying (91) on the left by E', we get

$$y_{Et} = \gamma y_{St} + \Lambda_S \eta_t + u_{St}.$$

The key moment is still $\mathbb{E}\left[\left(y_{Et} - \gamma y_{St}\right) z_t\right] = 0$, where the GIV is again $z_t = y_{St} - y_{Et}$ in the simple case where $\Lambda = \iota$ and $G\iota = \iota$ — see Section 8.2 for the general case. Hence, GIV generalizes to "spatial" models with common shocks (most spatial models do not have latent common shocks).

Another extension is a Bayesian interpretation of the GIVs. This way, we can interpret GIVs in a Bayesian framework – see Section 13. In the simplest cases with Gaussian shocks, the maximum likelihood estimate is our GIV – confirming its optimality properties. At the same time, the GIV doesn't actually use normality assumptions.

Furthermore, many econometric extensions might be useful, e.g. with stochastic volatility, and various dimensions of autocorrelations. We leave those extensions to future research.

7 Conclusion

In this paper we introduce the granular instrumental variables (GIVs): we remark that idiosyncratic shock offer a rich source of instruments, and we develop econometric procedures to extract them from aggregate shocks.

We provided some empirical applications. We plan to put on our web page a series of GIVs, and their control shocks η_t's. They might be useful for empirical work.

Many more are conceivable. Do firm-specific hiring and investment spill over to peer firms operating in the same product market? How much do constraints of financial intermediaries (e.g., broker dealers) matter for asset prices? How important are demand and supply shocks in oil markets? How much do international shock propagate ? How does a currency movement affect exports and borrowing (one could handle that via idiosyncratic demand shocks by large financial institutions for the currency in question)? We think that GIVs are likely to be helpful in a variety of settings beyond those studied in this paper.

References

Acemoglu, Daron, Ufuk Akcigit, and William Kerr, "Networks and the macroeconomy: An empirical exploration," *NBER Macroeconomics Annual*, 2016, *30* (1), 273–335.

_ , Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The network origins of aggregate fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Amiti, Mary and David E Weinstein, "How Much do Idiosyncratic Bank Shocks Affect Investment? Evidence from Matched Bank-Firm Loan Data," *Journal of Political Economy*, forthcoming.

_ , Oleg Itskhoki, and Jozef Konings, "International shocks and domestic prices: how large are strategic complementarities?," Staff Reports 771, Federal Reserve Bank of New York March 2016.

Auerbach, Alan J and Yuriy Gorodnichenko, "Measuring the output responses to fiscal policy," *American Economic Journal: Economic Policy*, 2012, *4* (2), 1–27.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, John Van Reenen et al., "The fall of the labor share and the rise of superstar firms," 2017.

Autor, David H, David Dorn, and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Azar, José, Ioana Marinescu, and Marshall I Steinbaum, "Labor market concentration," *NBER Working Paper No. 24147*, 2017.

Baqaee, David Rezza and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," NBER Working Papers 23145, National Bureau of Economic Research, Inc January 2018.

Bartik, Timothy J, "Who benefits from state and local economic development policies?," 1991.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile prices in market equilibrium," *Econometrica*, 1995, pp. 841–890.

Blasques, Francisco, Siem Jan Koopman, Andre Lucas, and Julia Schaumburg, "Spillover dynamics for systemic risk measurement using spatial financial time series models," *Journal of Econometrics*, 2016, *195* (2), 211 – 223.

Bloom, Nicholas, "The impact of uncertainty shocks," *Econometrica*, 2009, *77* (3), 623–685.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, "Identification of peer effects through social networks," *Journal of econometrics*, 2009, *150* (1), 41–55.

Carvalho, Vasco and Xavier Gabaix, "The great diversification and its undoing," *American Economic Review*, 2013, *103* (5), 1697–1727.

Chodorow-Reich, Gabriel, "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?," *NBER Working Paper No. 23577*, 2017.

Cochrane, John H., "The Dog That Did Not Bark: A Defense of Return Predictability," *The Review of Financial Studies*, 2008, *21* (4), 1533–1575.

Crouzet, Nicolas and Janice Eberly, "Understanding weak capital investment: The role of market concentration and intangibles," in "prepared for the Jackson Hole Economic Policy Symposium" 2018.

di Giovanni, Julian, Andrei A Levchenko, and Isabelle Mejean, "Foreign Shocks as Granular Fluctuations," 2018.

Fama, Eugene F and Kenneth R French, "Common risk factors in the returns on stocks and bonds," *Journal of financial economics*, 1993, *33* (1), 3–56.

Farhi, Emmanuel and Jean Tirole, "Deadly embrace: Sovereign and financial balance sheets doom loops," *The Review of Economic Studies*, 2017, *85* (3), 1781–1823.

Feenstra, Robert C. and David E. Weinstein, "Globalization, markups and U.S. welfare," *Journal of Political Economy*, 2017, *125* (4), 1040–1074.

Gabaix, Xavier, "The granular origins of aggregate fluctuations," *Econometrica*, 2011, *79* (3), 733–772.

_ , "A Behavioral New Keynesian Model," *NBER Working Paper No. 22954*, 2018.

Galí, Jordi and Mark Gertler, "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 1999, *44* (2), 195–222.

Gandhi, Priyank and Hanno Lustig, "Size Anomalies in U.S. Bank Stock Returns," *Journal of Finance*, 2015, *70* (2), 733–768.

Gaubert, Cecile and Oleg Itskhoki, "Granular comparative advantage," Technical Report, National Bureau of Economic Research 2018.

Giovanni, Julian Di and Andrei A Levchenko, "Country size, international trade, and aggregate fluctuations in granular economies," *Journal of Political Economy*, 2012, *120* (6), 1083–1132.

_ , _ , **and Isabelle Méjean**, "Firms, destinations, and aggregate fluctuations," *Econometrica*, 2014, *82* (4), 1303–1340.

Gutiérrez, Germán and Thomas Philippon, "Declining Competition and Investment in the US," Technical Report, National Bureau of Economic Research 2017.

_ **and** _ , "Declining Competition and Investment in the US," *NBER Working Paper No. 23583*, 2017.

Hamilton, James D., "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 2009, *40* (1 (Spring), 215–283.

He, Zhiguo, Bryan Kelly, and Asaf Manela, "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 2017, *126*, 1–35.

Ioannides, Yannis Menelaos, *From neighborhoods to nations: The economics of social interactions*, Princeton University Press, 2013.

Kelejian, Harry H and Ingmar R Prucha, "A generalized moments estimator for the autoregressive parameter in a spatial model," *International economic review*, 1999, *40* (2), 509–533.

Kilian, Lutz, "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market," *American Economic Review*, 2009, *99* (3), 1053–69.

Kline, Brendan and Elie Tamer, "Econometric Analysis of Models With Social Interactions," in Bryan Graham and Aureo De Paula, eds., *The Econometric Analysis of Network Data*, Academic Press, Forthcoming.

Koijen, Ralph SJ and Motohiro Yogo, "An equilibrium model of institutional demand and asset prices," *Journal of Political Economy*, forthcoming.

Kuersteiner, Guido M and Ingmar R Prucha, "Dynamic spatial panel models: Networks, common shocks, and sequential exogeneity," *arXiv preprint arXiv:1802.01755*, 2018.

Landier, Augustin, David Sraer, and David Thesmar, "Banking integration and house price co-movement," *Journal of Financial Economics*, 2017, *125* (1), 1–25.

Leary, Mark T and Michael R Roberts, "Do peer firms affect corporate financial policy?," *The Journal of Finance*, 2014, *69* (1), 139–178.

Lee, Youngki, Luís A Nunes Amaral, David Canning, Martin Meyer, and H Eugene Stanley, "Universal features in the growth dynamics of complex organizations," *Physical Review Letters*, 1998, *81* (15), 3275.

Loecker, Jan De and Jan Eeckhout, "The rise of market power and the macroeconomic implications," Technical Report, National Bureau of Economic Research 2017.

Long, John B and Charles I Plosser, "Real business cycles," *Journal of political Economy*, 1983, *91* (1), 39–69.

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Merton, Robert C, "On the pricing of corporate debt: The risk structure of interest rates," *The Journal of finance*, 1974, *29* (2), 449–470.

Mian, Atif and Amir Sufi, "What explains the 2007–2009 drop in employment?," *Econometrica*, 2014, *82* (6), 2197–2223.

Nakamura, Emi and Jon Steinsson, "Fiscal stimulus in a monetary union: Evidence from US regions," *American Economic Review*, 2014, *104* (3), 753–92.

_ **and** _ , "Identification in Macroeconomics," *Journal of Economic Perspectives*, August 2018, *32* (3), 59–86.

Ramey, Valerie A, "Macroeconomic shocks and their propagation," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 71–162.

Rasmussen, Carl and Christopher Williams, *Gaussian Processes for Machine Learning*, The MIT Press, 2005.

Romer, Christina D and David H Romer, "A new measure of monetary shocks: Derivation and implications," *American Economic Review*, 2004, *94* (4), 1055–1084.

Saiz, Albert, "The geographic determinants of housing supply," *The Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Sarto, Andres Pablo, "Recovering Macro Elasticities from Regional Data," 2018.

Shi, Wei and Lung fei Lee, "Spatial dynamic panel data models with interactive fixed effects," *Journal of Econometrics*, 2017, *197* (2), 323–347.

Sraer, David and David Thesmar, "A Sufficient Statistics Approach for Aggregating Firm-Level Experiments," NBER Working Papers 24208, National Bureau of Economic Research, Inc January 2018.

Stock, James H and Mark W Watson, "Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 415–525.

_ **and** _ , "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, 2018, *128* (610), 917–948.

_ **and Motohiro Yogo**, "Testing for weak instruments in linear IV regression," *NBER Technical Working Paper No. 284*, 2002.

Whelan, Karl, "A Guide To U.S. Chain Aggregated Nipa Data," *The Review of Income and Wealth*, 2002, *48* (2), 217–233.

Zou, Hui, Trevor Hastie, and Robert Tibshirani, "Sparse principal component analysis," *Journal of computational and graphical statistics*, 2006, *15* (2), 265–286.

8 Appendix: Complements and Extensions

8.1 Identifiability with factor models

Here we detail the procedure with parametric factor models, as in Section 3.2 Proof of Proposition 5. We next prove this proposition and more importantly show how to construct the GIVs.

The solution is, with $\lambda_S = S'\Lambda$ a $1 \times r$ vector, $M = \frac{1}{1-\lambda_S \alpha}$,

$$y_{St} = M\left(u_{St} + \lambda_S \eta_t + C_{St} m\right), \tag{92}$$

$$y_t = u_t + \Lambda\left[\alpha M\left(u_{St} + \lambda_S \eta_t + C_{St} m\right) + \eta_t\right] + C_t m. \tag{93}$$

Given a symmetric matrix W of size $N \times N$ (which, later, will optimally be $W = (V^u)^{-1}$, but we don't use that here yet) we define:[44]

$$Q = I - \Lambda\left(\Lambda' W \Lambda\right)^{-1} \Lambda' W, \tag{94}$$

so that

$$Q\Lambda = 0, \qquad Q'W\Lambda = 0, \qquad (I - Q)\, W^{-1} Q' = 0, \qquad Q^2 = Q. \tag{95}$$

Roughly, Q is the projection on the space orthogonal to the Λ, but with a dot product that depends on W. Hence, (93) implies:

$$Qy_t = Qu_t + QC_t m. \tag{96}$$

Defining, for a vector Y_t,

$$\check{Y}_t := QY_t, \tag{97}$$

we have

$$\check{y}_t = \check{u}_t + \check{C}_t m. \tag{98}$$

[44] For instance, in our basic example with uniform loading $\Lambda = \iota$, $Q = I - \iota E'$, where $E = \frac{W\iota}{\iota' W \iota}$.

The controls C_t^k are all assumed to have non-zero cross-sectional variation: this is what allows to identify their m. A variable that's an "aggregate control" without cross-sectional variation (e.g. a time fixed effect, or maybe the world price of oil if we study the macroeconomics of a small country not affecting it) will be classified as an F_t^f – it's in $\mathcal{F}^{\text{Exo},O}$ the set of observable, exogenous factors.

8.1.1 Estimating the variance-covariance matrix of the residuals, V^u

This is like in the iterative version of Section 19.9.1.[45]

First, we estimate m, using (98) Basically, we can estimate m by OLS. It's pretty easy, as we have $(N-r)\times T$ effective values to use (where r is the number of factors):

$$m^e = \mathbb{E}_T\left[\check{C}_t' W \check{C}_t\right]^{-1} \mathbb{E}_T\left[\check{C}_t' W \check{y}_t\right]. \tag{99}$$

Next, we estimate V^u In what follows, we assume that m was estimated without errors (XX to check the impact of that). We have $\check{u}_t := Q\left(y_t - C_t m^e\right) = Q u_t$, so that:

$$V^{\check{u}} = \mathbb{E}\left[\check{u}_t \check{u}_t'\right] = Q\mathbb{E}\left[u_t u_t'\right] Q. \tag{100}$$

We consider the case where we have a priori knowledge that V^u diagonal. Let us call $D^u = \left(V_{ii}^u\right)_{i=1\ldots N}$ and similarly for $D^{\check{u}} = \left(V_{ii}^{\check{u}}\right)_{i=1\ldots N}$ (so they are vectors of dimension N) and a new matrix $R_{ij} := \left(Q_{ij}^W\right)^2$. Then:

$$D_i^{\check{u}} := V_{ii}^{\check{u}} = \sum_j Q_{ij}^W V_{jj}^u Q_{ij}^W = \sum_j R_{ij} D_j,$$

i.e., $D^{\check{u}} = R D^u$ so we recover

$$D^{u,e} = R^{-1} D^{\check{u}}. \tag{101}$$

Parametric variant We can do a parametric variant. We parametrize $\ln \sigma_{u_i}^2 = \beta^\sigma x_i^\sigma$ for some vector of characteristics x_i^σ, e.g. log size or log volatility (and x_i^σ has 1 as the first component) — σ is just a superscript here, not an exponent. So, we estimate β^σ by regressing the log estimated variance from (101) on the characteristics:

$$\ln D_i^{u,e} = \beta^\sigma x_i^\sigma + \varepsilon_i,$$

and take the fitted values for the diagonal covariance matrix of the u_i's:

$$D_i^u = e^{\beta^\sigma x_i^\sigma}. \tag{102}$$

[45] Indeed, in the most basic case the only factor loading is $\Lambda = \iota$, so that $Q = I - \iota E'$ for some E with $\iota' E = 1$. For instance, $\check{u}_{it} = u_{it} - u_{Et}$.

Loop over W This gives a consistent estimate of D^u, for any W. Now, Bayesian considerations [11/16/18o] indicate that the optimum W is

$$W = (V^u)^{-1}. \tag{103}$$

So, we can loop: a good initial W is probably $1/var(y_i)$. This gives an estimate of D^u, and a new, better estimate of $W = diag(1/D_i^u)$. We can keep looping. We have consistently estimated the variance matrix of V^u.

We have identified V^u. We can now estimate γ.

8.1.2 Estimating multipliers α^f, M by GIV

We have identified W^u in the previous subsection. We show how to estimate the α^f and M. We set Q as in (94). Then, (96) gave[46]

$$Qy_t = Qu_t + QC_t m. \tag{105}$$

Let us define

$$\Gamma \coloneqq Q'S, \tag{106}$$

and define the GIV to be $z_{\Gamma t} \coloneqq \Gamma'(y_t - C_t m)$, which gives:

$$z_{\Gamma t} \;\coloneqq\; \Gamma'(y_t - C_t m) = \Gamma' u_t = u_{\Gamma t}. \tag{107}$$

This relation means that we identify $u_{\Gamma t}$ exactly, even though we estimate the η_t with errors.[47]

The GIV is possible if and only if

$$\Gamma \neq 0. \tag{108}$$

[46] As a motivation, let's consider the Bayesian estimate of η_t. Suppose for now that $\alpha = 0$, and $m = 0$ so that there is no loop. Then, the log likelihood L satisfies (assuming a flat prior on η_t)

$$-2L = (y_t - \Lambda\eta_t)' W^u (y_t - \Lambda\eta_t) - \ln|W^u|,$$

so that the MLE sets the derivative of this w.r.t. η_t to 0:

$$\Lambda' W^u (y_t - \Lambda\eta_t^e) = 0,$$

i.e.

$$\eta_t^e = (\Lambda' W^u \Lambda)^{-1} \Lambda' W^u y_t, \tag{104}$$

and the residual is

$$u_t^e = y_t - \Lambda\eta_t^e = Q^{W^u} y_t$$

for the value of (94) with $W = W^u$. (In the basic case where $\Lambda = \iota$,then $u_t^e = y_t - \iota y_{Et}$ with $E = \frac{W\iota}{\iota' W\iota}$). This suggests defining $z_t = S'u_t^e = S'Q^{W^u} y_t$, which is what we do for the general case with $\gamma = 0$.

[47] This may be surprising, but consider the following simple case to see how this is true: if $y_{it} = \lambda p_t + \eta_t + u_{it}$, then as $\Gamma'\iota = 0$, $y_{\Gamma t} = u_{\Gamma t}$. So we perfectly measure the $u_{\Gamma t}$. This relation with the Q generalizes that simple example with more complex factors.

This is exactly what motivated Assumption 1 mentioned in Proposition 5.[48]

We define $E = S - \Gamma = (I - Q')\, S$, so that

$$u_{St} = u_{\Gamma t} + u_{Et},$$

then (95) implies that $\mathbb{E}\left[u_{Et} u_{\Gamma t}\right] = S'\left(I - Q\right) V^u Q' S = 0$, so that we have the key relation:

$$\mathbb{E}\left[u_{Et} u_{\Gamma t}\right] = 0. \tag{109}$$

In addition, the value

$$z_{Ct} \coloneqq C_{St} m \tag{110}$$

is also a valid instrument. Hence, the following is an instrument:

$$z_t = z_{\Gamma t} + z_{Ct}, \tag{111}$$

and it is actually optimal, given (92). The reasoning is exactly as in Section 19.2. So, an instrument is the following optimal combination:

$$z_t = u_{\Gamma t} + C_{St} m. \tag{112}$$

Hence, (92) reads:

$$y_{St} = M\left(u_{St} + \lambda_S \eta_t + C_{St} m\right) = \left(u_{\Gamma t} + u_{Et} + \lambda_S \eta_t + C_{St} m\right) = M z_t + \varepsilon_t^{y_S} \tag{113}$$

for $\varepsilon_t^{y_S} = M\left(u_{Et} + \lambda_S \eta_t\right)$ uncorrelated with z_t. Hence, M will be consistently estimated by regressing y_{St}:

$$\mathbb{E}\left[\left(y_{St} - M z_t\right) z_t\right] = 0. \tag{114}$$

Likewise, for f an observable, endogenous control (i.e., one such that α^f need not be 0 a priori), we can regress:

$$F_t^f = \alpha^f M z_t + \varepsilon_t^{F^f}$$

with $\varepsilon_t^F = \eta_t + \alpha M \varepsilon_t^{y_S}$ so that we can estimate αM consistently by the OLS regression of F_t on z_t.

Recovering the cross-sectionally important exogenous factors F_t^f. If we have a cross-sectionally important factor η_t^f, we will want to control for it to gain statistical precision. To do so, we first want to extract the factor. For notational simplicity, we assume that we removed all the controls $C_t m$ (i.e., we replace y_t by $y_t - C_t m$).

[48] If $V^u S$ was spanned by the Λ, we could write $S = W \Lambda b$ for some vector b, and we'd have $Q'S = 0$, by 95. Conversely, if S is not spanned by the $V^u S$, the it is easy to check that $\Gamma \neq 0$.

We define the $r \times N$ matrix[49]

$$L_t \coloneqq \left(\Lambda_t' W \Lambda_t\right)^{-1} \Lambda_t' W, \tag{115}$$

so that

$$L_t \Lambda_t = I_r. \tag{116}$$

Next, using $Q\Lambda_t = 0$ in the factor structure (72) gives:

$$(I - Q)\, y_t = \Lambda_t F_t + (I - Q)\, u_t.$$

Premultiplying this by L_t gives:

$$L_t\, (I - Q)\, y_t = F_t + L_t\, (I - Q)\, u_t.$$

Hence, an estimate of the F_t is

$$F_t^e \coloneqq L_t\, (I - Q)\, y_t. \tag{117}$$

Indeed, we will have

$$F_t^e = F_t + \varepsilon_t^{F^e},$$

where $\varepsilon_t^{F^e} = -L_t\, (I - Q)\, u_t$ is a small error. In addition, this error is orthogonal to $z_{\Gamma t}$:[50]

$$\mathbb{E}\left[F_t^e z_{\Gamma t}\right] = 0, \tag{118}$$

so that the measurement error in the factors does not introduce a bias when estimating M.

Given this, decompose $F_t = \left(F_t^{\text{endo}}, F_t^{\text{exo}}\right)$ with endogenous factors (i.e., affected by u_{it}) and exogenous factors. We keep F_t^{exo} as use them as controls.

To get a feel for the situation, let us take the main example, where $\Lambda_{it} = (1, x_{it})$ with $x_{Et} = 0$ and $y_{it} = \gamma p_t + \varepsilon_t + x_i \eta_t + u_{it}$, and $W = I\sigma_u^{-2}$. So factors are: $F_t = \left(F_t^{\text{endo}}, F_t^{\text{exo}}\right) = (\gamma p_t + \varepsilon_t, \eta_t)$, we have $I - Q = \Lambda L$, so $L\,(I - Q) = L\Lambda L = L$, hence $L y_t = \left(y_{Et}, \frac{x_t' y_t}{\|x_t\|^2}\right)$, so that the error is

$$F_t^e - F_t = \varepsilon_t^{F^e} = -L\,(I - Q)\, u_t = \left(u_{Et}, \frac{x_t' u_t}{\|x_t\|^2}\right), \tag{119}$$

and the standard deviation of its components are $\frac{\sigma_u}{\sqrt{N}}, \frac{\sigma_u}{\sqrt{N}} \frac{1}{\sigma_x}$. The factor analysis recovers (up to that error) $F_t = \left(F_t^{\text{endo}}, F_t^{\text{exo}}\right) = (\gamma p_t + \varepsilon_t, \eta_t)$, so it recovers η_t, but not ε_t. We can use that η_t as a control in the regression.

[49] We might also take $L_t \coloneqq \left(\Lambda_t' \Lambda_t\right)^{-1} \Lambda_t'$.

[50] Indeed, (107) gives

$$-\mathbb{E}\left[F_t^e z_t\right] \quad = \quad L\,(I - Q)\, V^u Q' S = L\,(I - Q)\, W^{-1} Q' S = 0$$

using (95) $(I - Q)\, W^{-1} Q' = 0$.

Regression with controls If we have observable, exogenous factors F_t^g with $g \in \mathcal{F}^{\text{Exo},O}$, then for efficiency it is good to include them as controls, and regress by OLS:

$$y_{St} = Mz_t + \sum_{g \in \mathcal{F}^{\text{Exo},O}} \beta_g^{yS} F_t^g + \tilde{\varepsilon}_t^{yS}, \tag{120}$$

$$F_t^f = \alpha^f M z_t + \sum_{g \in \mathcal{F}^{\text{Exo},O}} \beta_g^{yS} F_t^g + \tilde{\varepsilon}_t^F. \tag{121}$$

Hence, there are three types of "controls": (i) First, the controls C_{it}^k that have cross-sectional variation, so that their impact m^k can be estimated from the cross-section. They can be used as instruments in $z_{Ct} \coloneqq \sum_k C_{St}^k m_k$. (ii) Second, the controls that impact the whole system without cross-sectional differences, so that their impact cannot be estimated in the cross-section, F_t^g (with $g \in \mathcal{F}^{\text{Exo},O}$). (iii) Hidden factors that are recovered from the factor analysis. The last two types can only be used as pure controls in the above regressions, but not as instruments.

In conclusion, with a factor model (with known factor loadings Λ, but unknown factor covariance matrix V^η), we have identified V^u, and gotten a GIV, which gave M, αM.

Even though all worked with finite N (but as always, $T \to \infty$), and we don't consistently estimate η_t, we still have a consistent estimator for the GIV.

8.2 When the influence matrix is not proportional to size

8.2.1 Position of the problem

Suppose a model

$$y_{it} = \gamma \sum_j G_{ij} y_{jt} + \lambda_i \eta_t + u_{it}, \tag{122}$$

i.e.

$$y_t = \gamma G y_t + \Lambda \eta_t + u_t, \tag{123}$$

with a given "influence" matrix G. For instance, if we have an "industrial similarity" matrix H_{ij}, we might set

$$G_{ij} = \frac{H_{ij} S_j}{\sum_k H_{ik} S_k}.$$

In our basic "reflection problem", $G = \iota S'$.

We'd like to identify γ.

With $V = \mathbb{E}[u_t u_t']$, we define $\tilde{E} = \frac{V^{-1}\iota}{\iota V^{-1}\iota}$, and

$$S \coloneqq G' \tilde{E}, \tag{124}$$

which is the analogue of "size" in our simpler setup where $G = \iota S'$.

8.2.2 A simple approach, when the loading on common shocks if uniform

In this subsection we assume that

$$G\iota = \iota, \tag{125}$$

which is satisfied in many examples. Consider some vector $\mathcal{E}$, and define:

$$z_t \coloneqq \mathcal{E}' (G - I) y_t. \tag{126}$$

Then, we have the key relation:

$$\mathbb{E}\left[z_t \left(y_{\tilde{E}t}\text{-}\gamma y_{St}\right)\right] = 0, \tag{127}$$

which allows to identify γ by $\gamma = \frac{\mathbb{E}\left[z_t y_{\tilde{E}t}\right]}{\left[z_t y_{St}\right]}$.

Relation (127) works for any z_t of the type (126). We can show that the optimum choice (in the sense of minimizing the asymptotic error) is $\mathcal{E} = \tilde{E}$.[51] Then, the GIV is again:

$$z_t = y_{St} - y_{\tilde{E}t}. \tag{128}$$

Derivation of (127) Indeed,

$$y_{\tilde{E}t} - \gamma y_{St} = \tilde{E}' (I - \gamma G) y_t = \tilde{E}' (u_t + \eta_t \iota).$$

Given $G\iota = \iota$, and $(G - I)$ and $(I - \gamma G)^{-1}$ commute, we have the useful relation:

$$(G - I) (I - \gamma G)^{-1} \iota = 0. \tag{129}$$

As a result,

$$z_t = \mathcal{E}' (G - I) y_t = \mathcal{E}' (G - I) (I - \gamma G)^{-1} (u_t + \eta_t \iota) = \mathcal{E}' (G - I) (I - \gamma G)^{-1} u_t.$$

Hence,

$$\begin{aligned} a &\coloneqq \mathbb{E}\left[z_t \left(y_{\tilde{E}t} - \gamma y_{St}\right)\right] = \mathbb{E}\left[\mathcal{E}' (G - I) (I - \gamma G)^{-1} u_t (u_t + \eta_t \iota)' \tilde{E}\right] \\ &= \mathcal{E}' (G - I) (I - \gamma G)^{-1} V \tilde{E} = \mathcal{E}' (G - I) (I - \gamma G)^{-1} \frac{\iota}{\iota V^{-1} \iota} = 0. \end{aligned}$$

because of (129). □

[51] Indeed, one can show that this is the MLE estimator [10/27/18 notes].

8.2.3 When the loading on common shocks is not necessarily uniform

We now study the more general case where:

$$y_t = \gamma G y_t + \Lambda \eta_t + u_t, \tag{130}$$

where the factor loading Λ (an $N \times r$ matrix) is not necessarily equal to ι (but we keep imposing that the Λ spans ι, i.e. there is a q such that $\iota = \Lambda q$). As before, η_t is a low-dimensional vector of factors. However, we do not assume anymore that $G\iota = \iota$.

First, we suppose that we have a first estimate of γ, which we call if γ^e. We will later iterate on it. Then, we form:

$$\tilde{y}_t(\gamma^e) = (I - \gamma^e G) y_t. \tag{131}$$

If $\gamma^e = \gamma$, then

$$\tilde{y}_t(\gamma) = \Lambda \eta_t + u_t. \tag{132}$$

Hence, we run a factor analysis on $\tilde{y}_t(\gamma^e)$, which recovers Λ and $W = (V^u)^{-1}$. We introduce Q as in (94) so that $Q\Lambda = 0$ and set

$$\check{u}^e = Q \tilde{y}_t(\gamma^e),$$

so that at $\gamma^e = \gamma$,

$$\check{u}^e_t = Q u_t. \tag{133}$$

Then, we define (with $E = \frac{W\iota}{\iota' W \iota}$), with $S \coloneqq G'E$,[52][53]

$$\begin{aligned} z_t &= S' Q (1 - \gamma^e G)^{-1} \check{u}^e_t & (134)\\ &= S' Q (1 - \gamma^e G)^{-1} Q (I - \gamma^e G) y_t. & (135) \end{aligned}$$

Our key moment is as before, equation (127):[54]

$$\mathbb{E}\left[z_t \left(y_{\tilde{E}t} - \gamma y_{St}\right)\right] = 0. \tag{137}$$

[52] A variant is:

$$z_t = S' (1 - \gamma^e G)^{-1} \check{u}^e_t = S' (1 - \gamma^e G)^{-1} Q (I - \gamma^e G) y_t.$$

The advantage of formulation (127) is that in the simple case of the previous subsection (with $\Lambda = \iota$), then it recovers the estimator of that subsection.

[53] The proof shows that this choice works. This particular choice is heuristically motivated by the analogy with (107) and (124)), and the fact that when $\eta_t = 0$, $y_t = (I - \gamma G)^{-1} u_t$, so that $y_{St} = S' y_t = \gamma S' (1 - \gamma^e G)^{-1} u_t$. Hence in some loose sense z_t is a good idiosyncratic-based approximation of y_{St}.

[54] Here is the proof. At the right estimator $\gamma = \gamma^e$,

$$z_t = S' Q (1 - \gamma^e G)^{-1} Q u_t = c' Q u_t, \tag{136}$$

with $c' \coloneqq S' Q (1 - \gamma^e G)^{-1}$. We also have

$$y_{\tilde{E}t} - \gamma y_{St} = E' (I - \gamma G) y_t = E' (\Lambda \eta_t + u_t).$$

This yields an estimate of γ. Hence, we simply replace the definition of the GIV (126) by (135). In fact, we can show that when $\Lambda = \iota$, the estimator (135) is equal to the estimator (126). In that sense it is its natural generalization.[55]

We note that we have a fixed point: an initial γ^e gives an estimate of γ; that's then the new estimate γ^e, and we re-iterate the process, until convergence.

Suppose that instead we use[56]

$$Z_t = S' \left(1 - \gamma^e G\right)^{-1} \breve{u}_t^e. \tag{138}$$

Suppose that $\gamma^e = \gamma$. Then we can write:[57]

$$y_{St} = Z_t + \varepsilon_t, \qquad \mathbb{E}\left[\varepsilon_t Z_t\right] = 0, \tag{139}$$

This implies that

$$\mathbb{E}\left[\left(y_{\tilde{E}t} - \gamma y_{St}\right) z_t\right] \quad = \quad \mathbb{E}\left[E' u_t u_t' Q' c\right] = E' V Q' c = \frac{1}{\iota' W \iota} \iota' Q' c = 0,$$

as $Q\iota = 0$.

[55] In the case $\Lambda = \iota$, then $Q = I - \iota E'$, so

$$E'GQ = E'G\left(I - \iota E'\right) = E'G - E' = E'\left(G - I\right),$$

so that the estimator in 135 can be written:

$$z_t = E'\left(G - I\right) y_t = E'GQy_t.$$

On the other hand,

$$\begin{aligned} A &\coloneqq GQ\left(1 - \gamma^e G\right)^{-1} Q\left(I - \gamma^e G\right) = GQ\left(1 - \gamma^e G\right)^{-1}\left(I - \iota E'\right)\left(I - \gamma^e G\right) \\ &= GQ\left(1 - \gamma^e G\right)^{-1}\left(I - \gamma^e G\right) - GQ\left(1 - \gamma^e G\right)^{-1} \iota E'\left(I - \gamma^e G\right) \\ &= GQ - 0 \text{ as } G\iota - \iota \text{ and } Q\iota = 0 \\ &= GQ, \end{aligned}$$

so that the estimator in (135) can be written:

$$z_t = E'GQ\left(1 - \gamma^e G\right)^{-1} Q\left(I - \gamma^e G\right) y_t = E' A y_t = E'GQy_t = E'\left(G - I\right) y_t.$$

Hence, when $\Lambda = \iota$, the estimators in (126) and (135) are identical.

[56] Note that the γ^e in the definition of Z_t need note be the same γ^e used above to construct the $\breve{u}_t^e$; i.e., we could have $Z_t = S'\left(1 - \gamma^{e,2} G\right)^{-1} \breve{u}_t^e$ for some other $\gamma^{e,2}$. There is still a fixed point though, and in the limit the estimated γ in (140) should also be equal to the γ^e and $\gamma^{e,2}$.

[57] Here is the proof. We saw that $\breve{u}_t^e = Qu_t$, so,

$$Z_t = S'\left(I - \gamma^e G\right)^{-1} Q u_t,$$

while, $a' \coloneqq S'\left(I - \gamma^e G\right)^{-1}$,

$$y_{St} = S'\left(I - \gamma^e G\right)^{-1}\left(Qu_t + \left(I - Q\right) u_t + \Lambda \eta_t\right) = Z_t + a'\left(I - Q\right) u_t + a' \Lambda \eta_t.$$

From (95), we have $\mathbb{E}\left[a'\left(I - Q\right) u_t z_t\right] = 0$. Hence we have $\mathbb{E}\left[\varepsilon_t Z_t\right] = 0$.

for $\varepsilon_t = S'\left(I - \gamma^e G\right)^{-1}\left(\left(I - Q\right)u_t + \Lambda\eta_t\right)$. Also, we have:

$$y_{Et} = \gamma y_{St} + E'\Lambda\eta_t + u_{Et}.$$

Hence, we can estimate γ by OLS:

$$y_{Et} = \gamma Z_t + \beta'\eta_t + \varepsilon_t^{y_E}. \tag{140}$$

This consistently estimates γ.

Calling z_t the GIV in the basic models (with $G = \iota' S$), $Z_t = \frac{z_t}{1-\gamma^e}$.

8.3 Link between our initial examples and the general framework

Let us make the link with our initial examples and the general setup of Section 3.1.

The basic example of Section 2.2 was:

$$y_{it} = \eta_t + u_{it}, \qquad p_t = \alpha y_{St} + \varepsilon_t.$$

The factors are p_t and η_t:

$$F_t^1 = p_t, \quad \alpha^1 = \alpha, \quad \lambda^1 = 0, \quad \eta_t^1 = \varepsilon_t,$$

$$F_t^2 = \eta_t, \quad \alpha^2 = 0, \quad \lambda^2 = 1, \quad \eta_t^2 = \eta_t.$$

Factor F_t^1 is endogenous, factor F_t^2 is exogenous. They are both observable (via $y_{Et} = \eta_t + O\left(\frac{1}{\sqrt{N}}\right)$). So, $\mathcal{C}_t = \eta_t$. The multiplier is $M = 1$.

Let us next make the link with supply-and-demand example of Section 2.3, which was:

$$y_{it} = \lambda p_t + \eta_t + u_{it}, \qquad p_t = \frac{y_{St}}{\mu} - \frac{\varepsilon_t}{\mu}.$$

The factors are also p_t and η_t:

$$F_t^1 = p_t, \quad \alpha^1 = \frac{1}{\mu}, \quad \lambda^1 = \lambda, \quad \eta_t^1 = -\frac{\varepsilon_t}{\mu},$$

$$F_t^2 = \eta_t, \quad \alpha^2 = 0, \quad \lambda^2 = 1, \quad \eta_t^2 = \eta_t.$$

Here F_t^1 is exogenous and hidden, while F_t^2 is endogenous and observable. So, $\mathcal{C}_t$ is empty.[58] The multiplier is $M = \frac{1}{1-\alpha^1\lambda^1} = \frac{\mu}{\mu-\lambda}$, as was estimated in (47).

[58] This is why we could do only plain OLS regression in (47), without any controls like y_{Et}, unlike in the very simple initial example of Section 2.2.

8.4 Derivations of the passage from regional to macro impacts

Here we give the microfoundations for the results in Section 5.

The aggregate economy Let us start with the aggregate economy, to get a feel for the basic ingredients. They are drawn from Gabaix (2018), (in particular, the simplified model of Section 9), where the detailed microfoundations are laid out.

The consumer has utility $\sum_{t=0}^{\infty} \beta^t \ln c_t$. For simplicity, we normalize all initial quantities and prices to be 1, and denote by hats the fractional deviations from the steady state. The gross interest rate is $R = \frac{1}{\beta}$.

With μ the marginal propensity to consume out of current income, we have:

$$\hat{c}_t = \mu \hat{y}_t + w_t,$$

where w_t captures a perceived present value of future income.

To familiarize ourselves with this environment, suppose that the government consumes G_0, financed by taxes at time 1, of which only a fraction m_G is anticipated. We have $w_0 = -\mu m_G G_0$, so that:

$$\hat{c}_0 = \mu \left(\hat{y}_0 - m_G G_0 \right). \tag{141}$$

As in general equilibrium $\hat{y}_0 = \hat{c}_0 + G_0$, we have:

$$\hat{y}_0 - G_0 = \mu \left(\hat{y}_0 - m_G G_0 \right), \tag{142}$$

which yields a "Keynesian cross"

$$\hat{y}_0 = \frac{\left(1 - \mu m_G\right) G_0}{1 - \mu}. \tag{143}$$

We take the benchmark where consumers do not anticipate future taxes – they are non-Ricardian, $m_G = 0$. Then, the impact of government spending is $\frac{1}{1-\mu}$.

This machinery in place, we move to the more complex case with several regions.

The economy with N regions There are N regions. The representative consumer in region i has utility

$$u\left(c_{it}^{NT}, c_{it}^{T}\right) = \alpha \ln c_{it}^{NT} + (1 - \alpha) \ln c_{it}^{T}.$$

Then, the Cobb-Douglas form with constant prices implies:

$$\hat{c}_{it}^{N} = \hat{c}_{it}^{NT} = \hat{c}_{it}.$$

Eq. (141) can be summarized as: $\hat{c}_{it} = \mu \hat{y}_{it} + v_{it}$, where v_{it} is a disturbance, e.g. about expectation about future income, including housing wealth. We allow for the possibility of a richer setup with:

$$\hat{c}_{it} = \mu \hat{y}_{it} + \kappa \hat{y}_{St} + v_{it}, \tag{144}$$

where the term $\kappa \hat{y}_{St}$ could be some other "expectational shock", e.g. via perhaps a learning channel, or some other sort of general equilibrium channel. It is not necessary to add it, but the fact that our propositions work even with that setup shows that they are robust to many variants.

Proof of Proposition 6 There is no government spending, so non-tradable GDP is made entirely of consumption. So, total income is:

$$\hat{y}_{it} = \alpha \hat{c}_{it} + (1 - \alpha)\, \hat{c}_{St}, \tag{145}$$

because region i sells a fraction α of its output locally, and $1 - \alpha$ in the whole country. Using (144) gives:

$$\hat{y}_{it} = b \hat{y}_{it} + (B - b)\, \hat{y}_{St} + \alpha v_{it} + (1 - \alpha)\, v_{St}. \tag{146}$$

with

$$b = \alpha \mu, \quad B = \mu + \alpha \kappa, \tag{147}$$

Aggregating this across all regions, $\hat{y}_{St} = B \hat{y}_{St} + v_{St}$, so

$$\hat{y}_{St} = \frac{1}{1 - B} v_{St}.$$

Hence, (146) gives, for some k:

$$\hat{y}_{it} = \frac{\alpha v_{it}}{1 - b} + k \hat{y}_{St} = u_{it} + k \hat{y}_{St},$$

where

$$u_{it} \coloneqq \frac{\alpha}{1 - b} v_{it} \tag{148}$$

is the local shock "after a first round of local multiplier $\frac{1}{1-b}$"; it is what we observe in practice. We relate it here to the more primitive (and not directly observed) v_{it}. Hence we have:

$$\hat{y}_{St} = \frac{1}{1 - B} v_{St} = \frac{1}{1 - B} \frac{1 - b}{\alpha} u_{St} = M u_{St},$$

with

$$M = \frac{1 - b}{1 - B} \frac{1}{\alpha}, \tag{149}$$

$$\hat{y}_{St} = M u_{St}. \tag{150}$$

Next, an econometrician in this economy, doing a GIV, will estimate this $M = \frac{d\hat{y}_{St}}{du_{St}}$.

Now, imagine that there are (say, exogenous at this stage) housing shocks h_{it}, that imply a local shock to total local employment

$$\hat{y}_{it} = \psi^{\ell} h_{it}, \tag{151}$$

as measured by Mian and Sufi (2014). This is equivalent to parametrizing

$$u_{it} = \psi^{\ell} h_{it}.$$

Hence, the impact on the aggregate is

$$\hat{y}_{St} = M u_{St} = M \psi^{\ell} h_{St} = \psi h_{St},$$

with $\psi = M\psi^{\ell}$. This proves the proposition.

Proof of Proposition 7 The steps are very similar to the previous proof, though complicated by government spending. We say that if the government spends G_{it} in region i, a fraction α^G is spent on tradables, and $1 - \alpha^G$ on nontradables. Hence, (145) becomes:

$$\hat{y}_{it} = \alpha \hat{c}_{it} + (1 - \alpha)\, \hat{c}_{St} + \alpha^G G_{it} + \left(1 - \alpha^G\right) G_{St}, \tag{152}$$

as income comes from local and global consumption by individuals and the government. Using (144) and for simplicity setting expectational shocks to $v_{it} = 0$, we have: $\hat{c}_{it} = \mu \hat{y}_{it} + \kappa \hat{y}_{St} + v_{it}$, $\hat{c}_{St} = (\mu + \kappa)\, \hat{y}_{St}$ hence, using the same values b, B as in (147)

$$\hat{y}_{it} = b \hat{y}_{it} + (B - b)\, \hat{y}_{St} + \alpha^G G_{it} + \left(1 - \alpha^G\right) G_{St}. \tag{153}$$

Aggregating, we get:

$$\hat{y}_{St} = B \hat{y}_{St} + G_{St}, \tag{154}$$

i.e.

$$\hat{y}_{St} = \frac{1}{1 - B} G_{St}, \tag{155}$$

so that the global impact is

$$\psi = \frac{\partial \hat{y}_{it}}{\partial G_{it}} = \frac{1}{1 - B}.$$

So, (153) indicates that the local impact is

$$\hat{y}_{it} = \frac{\alpha^G G_{it}}{1 - b} + k G_S$$

for some k. Hence, the local impact is

$$\psi^{\ell} = \frac{\partial \hat{y}_{it}}{\partial G_{it}} = \frac{\alpha^G}{1-b}.$$

Hence, inspecting (149), we obtain:

$$\psi = M\psi^{\ell}\frac{\alpha}{\alpha^G}, \tag{156}$$

as announced.